Parallel Enhancement
by High Performance Computing
+ GPU Parallel

Convergence and Collaboration

TAESUNG S&E, INC.

목 차

- HPC 기술 동향
- ANSYS HPC 라이선스
- ANSYS Mechanical의 HPC와 GPU 성능
- ANSYS CFD의 HPC와 GPU 성능

HPC 기술 동향

산업기술동향
HPC를 위한 요구

짧은 시간에 제품 성능 향상 요구
- 더 많은 설계 변경 고려
- 최적의 설계 조건 검색
- 주어진 조건 범위에서 성능 확보

➢ HPC는 강건 설계를 통해 더 향상된 제품의 무결성을 보증하는 핵심 요소

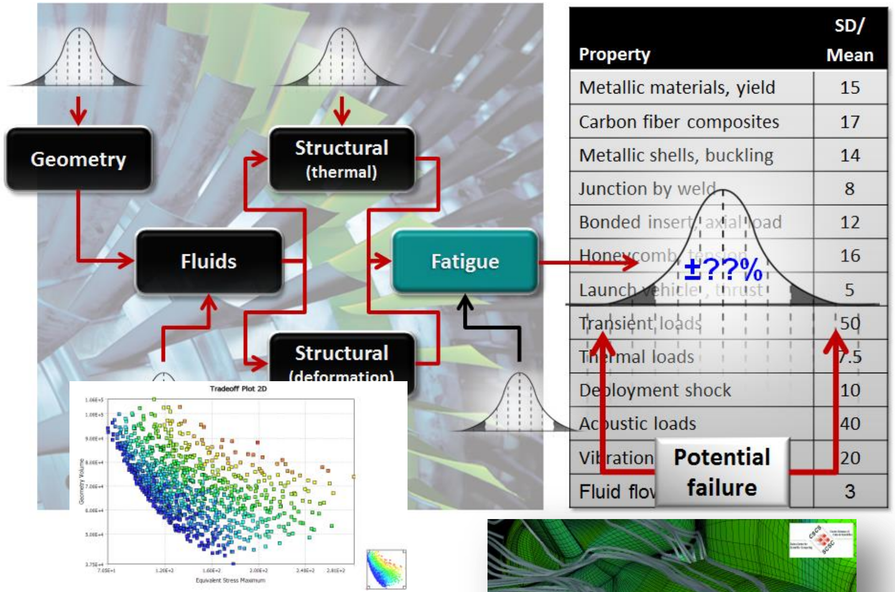

Property	SD/ Mean
Metallic materials, yield	15
Carbon fiber composites	17
Metallic shells, buckling	14
Junction by weld	8
Bonded insert, axial load	12
Honeycomb, tension	16
Launch vehicle, thrust	5
Transient loads	50
Thermal loads	7.5
Deployment shock	10
Acoustic loads	40
Vibration	20
Fluid flo	3

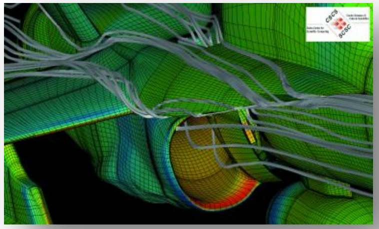

제품의 복잡성 증대
- 더 크고, 더 자세한 모델 평가
- 더 복잡한 물리계 고려
- 단일 요소에서 시스템 혁신까지 고려

➢ HPC는 대용량의 해석 고찰을 위한 핵심 요소

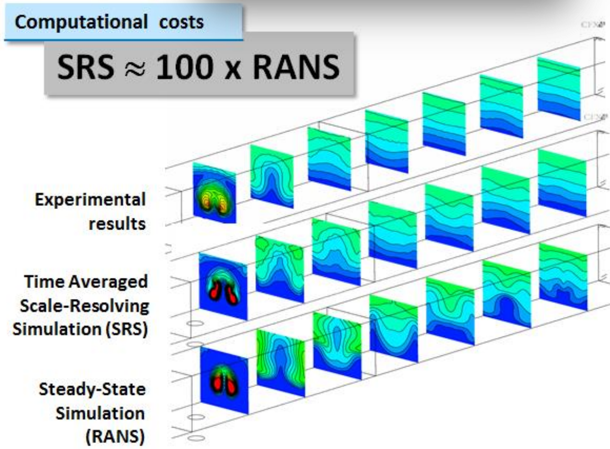

HPC를 위한 요구

엔지니어링 및 디자인 팀과 더 많은 협업 수행

- 더 많은 시뮬레이션 작업을 지원하기 위해 **HPC** 비율 증대
- 원격 접속을 통한 중앙 집중화된 **HPC** 인프라 구축(혹은 클라우드)
- 협업 작업을 통한 엔지니어링 데이터 관리 필요

➢**HPC**는 엔지니어링 생산성을 증대하는 핵심 요소

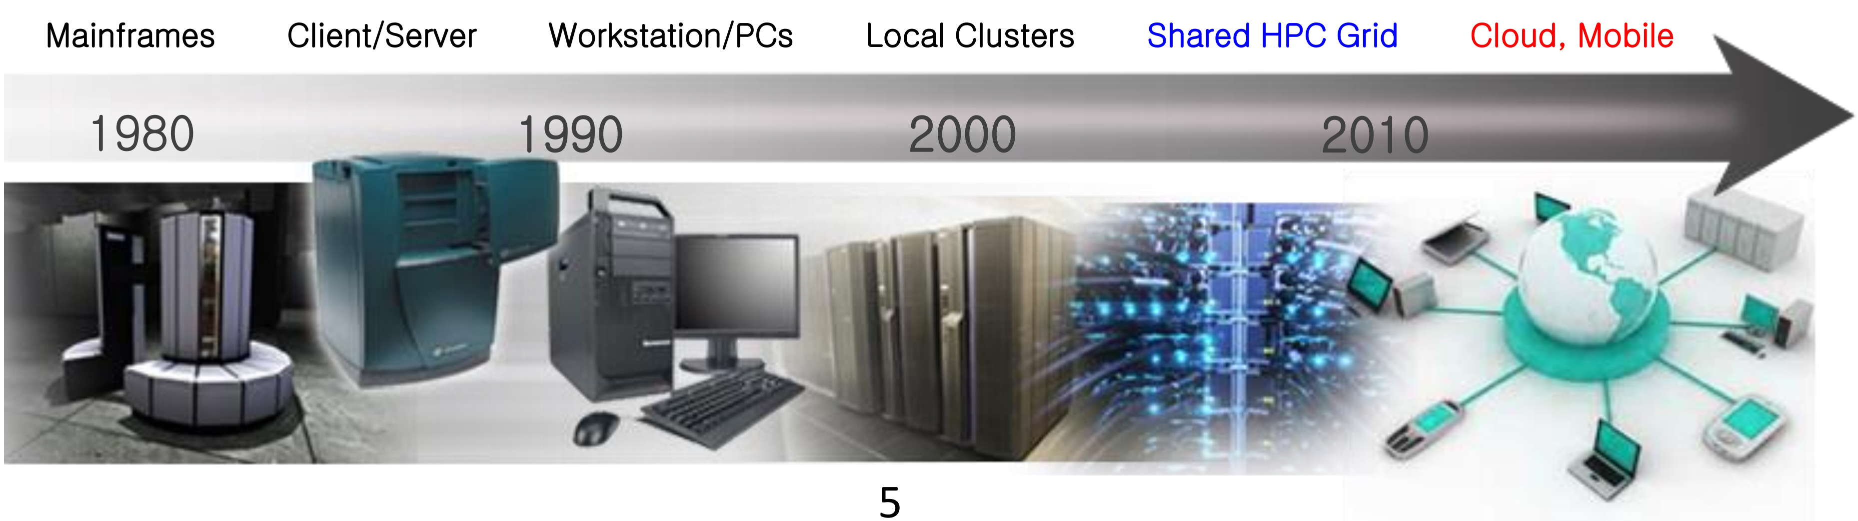

고성능 계산의 비례 증가

HPC 소프트웨어 개발의 빠른 속도

Single-Core 시대

~에 의해 가능 :
- √ **Moore**의 법칙
- √ 전압 스케일링
- √ 마이크로 아키텍쳐

~에 의해 구속:
- X 전력
- X 복잡성

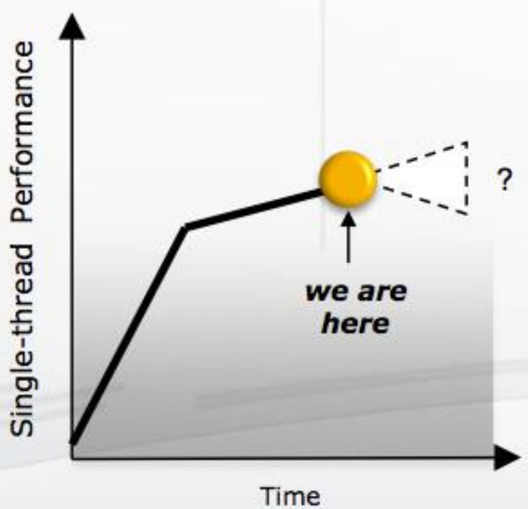

Multi-Core 시대

~에 의해 가능 :
- √ **Moore**의 법칙
- √ 처리량을 위한 갈망
- √ **SMP** 아키텍쳐의 **20**년

~에 의해 구속:
- X 전력
- X 병렬 소프트웨어의 이용성
- X 확장성

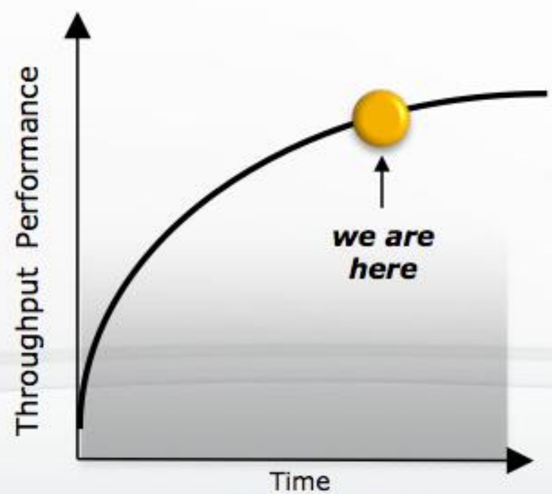

Heterogeneous(이종)-Core 시대

~에 의해 가능 :
- √ **Moore**의 법칙
- √ 풍부한 데이터 병렬 처리
- √ 효율적 전력 소모 **GPU**

한시적으로 ~에 의해 구속:
- X 프로그래밍 모델
- X 커뮤니케이션 오버헤드

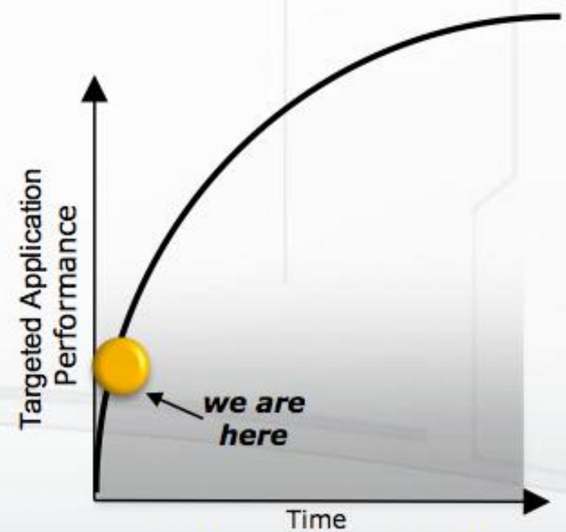

Source: AnandTech

ANSYS와 IT업체의 파트너쉽
성능/플랫폼 지원

IT 산업 리더와의 파트너쉽, 최적화된 HPC 성능 확보, 미래를 위한 로드맵, 밀착 지원

- ANSYS와 Intel – “Sandy Bridge” 프로세서로 60% 성능 향상; R&D는 Intel’s Xeon Phi 에 초점
- ANSYS와 NVIDIA – ANSYS Mechanical의 GPU 가속; 전체 포트폴리오에 따른 매우 활동적인 R&D 연대
- ANSYS와 Cray – 1000코어 이상의 극한 성능 향상 지원

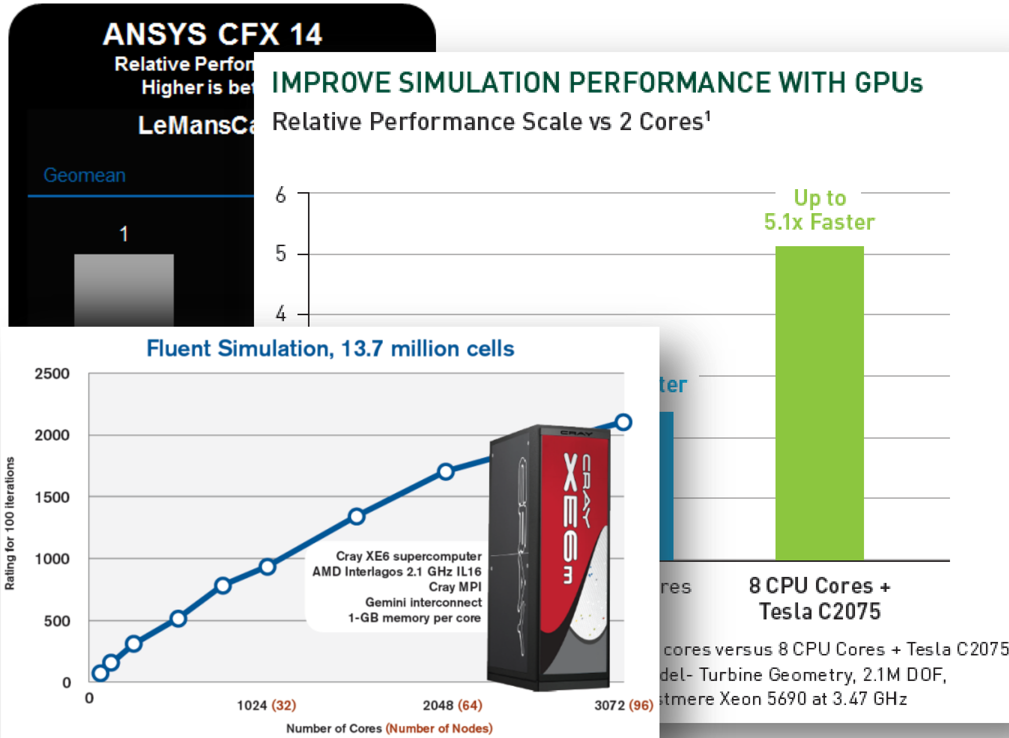

ANSYS HPC 성능은 라이선스, 하드웨어, 인력의 활용을 최적화한다.

ANSYS HPC 라이선스

HPC 라이센스

Scalable licensing

- ANSYS HPC (per-process)
- ANSYS HPC Pack
  - "Virtually Unlimited" parallel on single job
  - Best solution for extreme scaling needs
  - Parallel enabled increases quickly as packs added
- ANSYS HPC Workgroup
  - Volume access to HPC for many users running 'everyday' HPC jobs
- ANSYS HPC Enterprise
  - Similar to HPC Workgroup but deploy and use anywhere in the world

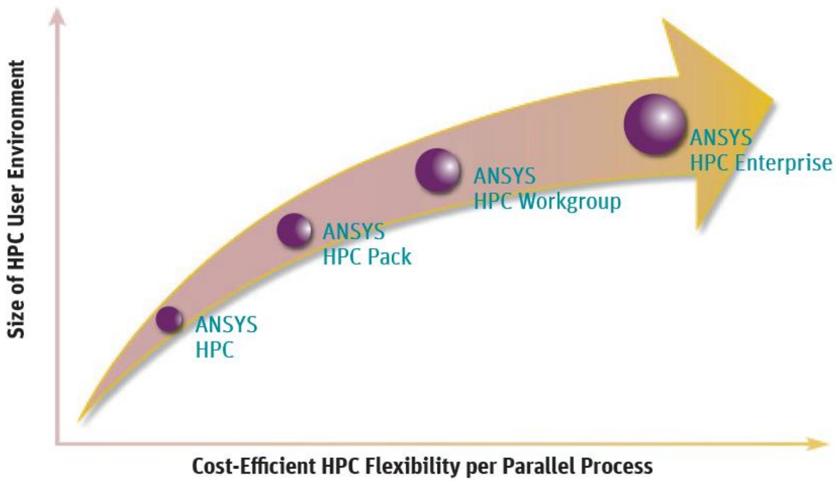

Single HPC solution for FEA/CFD/FSI and any level of fidelity

ANSYS HPC Pack

- **ANSYS HPC Pack**은 대용량의 해석에 대한 고찰이 가능
- **ANSYS HPC Pack**은 모든 물리계와 모든 사용자가 공용으로 대용량 해석을 위해 사용 가능
- **ANSYS HPC Pack**은 **GPU** 계산을 추가적으로 지원
- **HPC Pack** 라이선스는 추가될 때 마다 **Core+GPU** 사용 가능 수가 **4**배로 확장

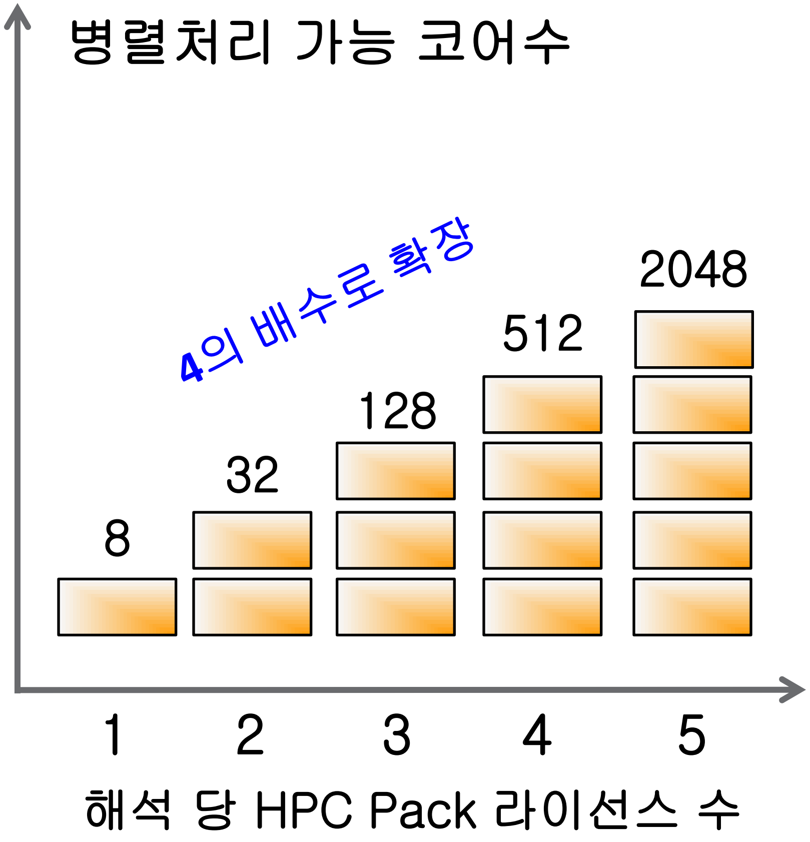

ANSYS HPC Parametric Pack

Explore Your Parametric Designs Faster, More Cost Effectively

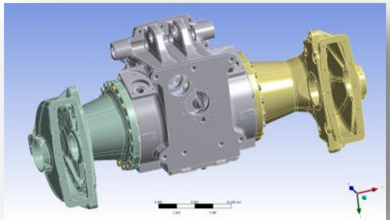

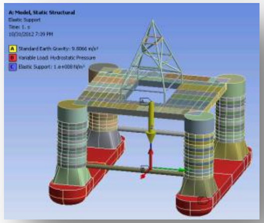

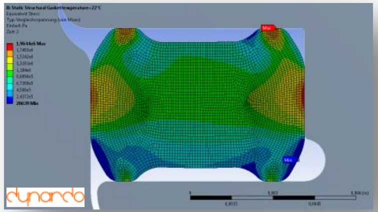

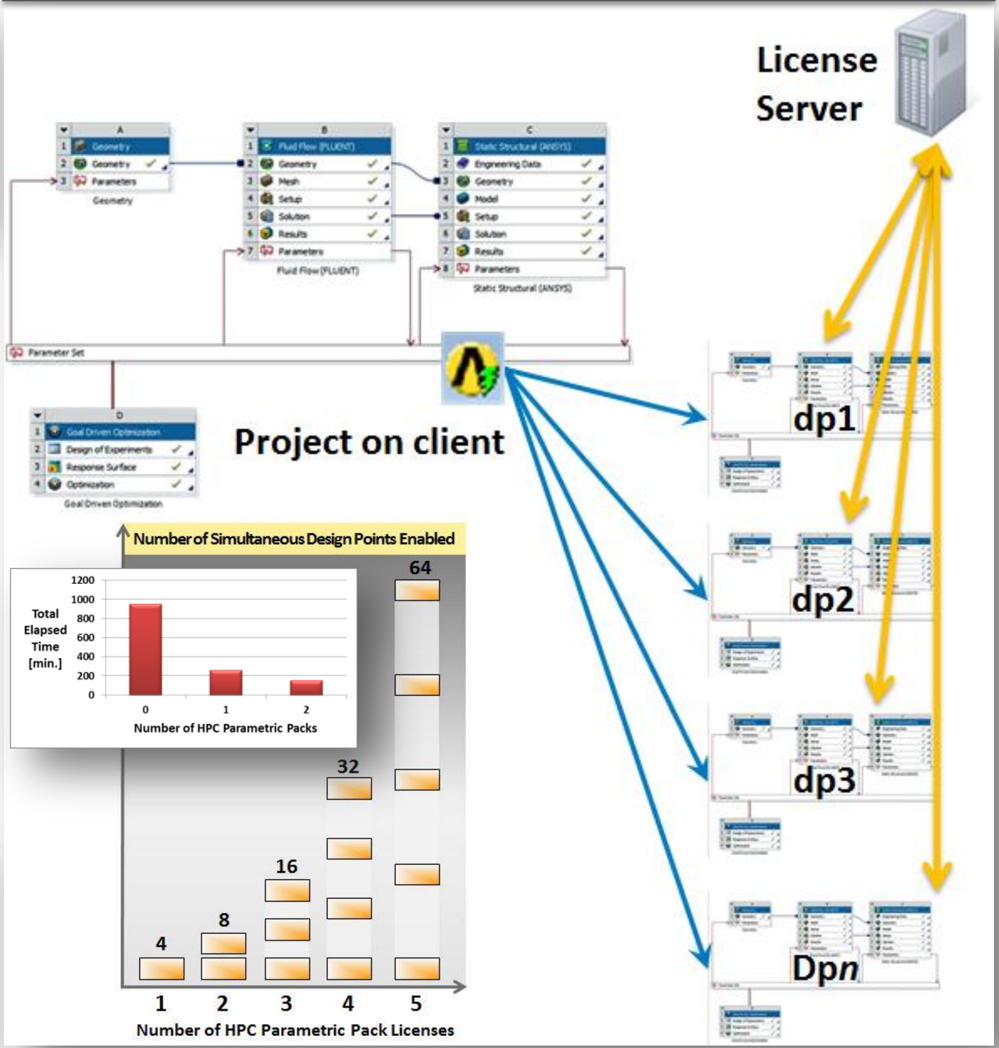

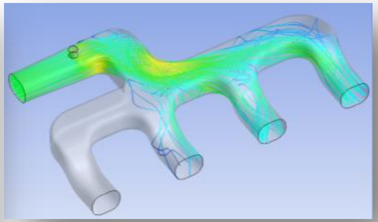

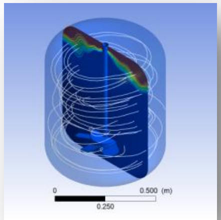

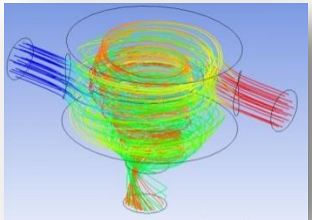

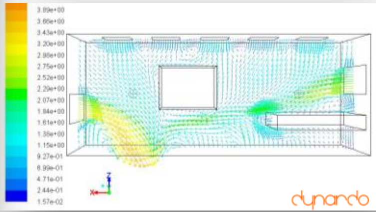

ANSYS HPC Parametric Pack

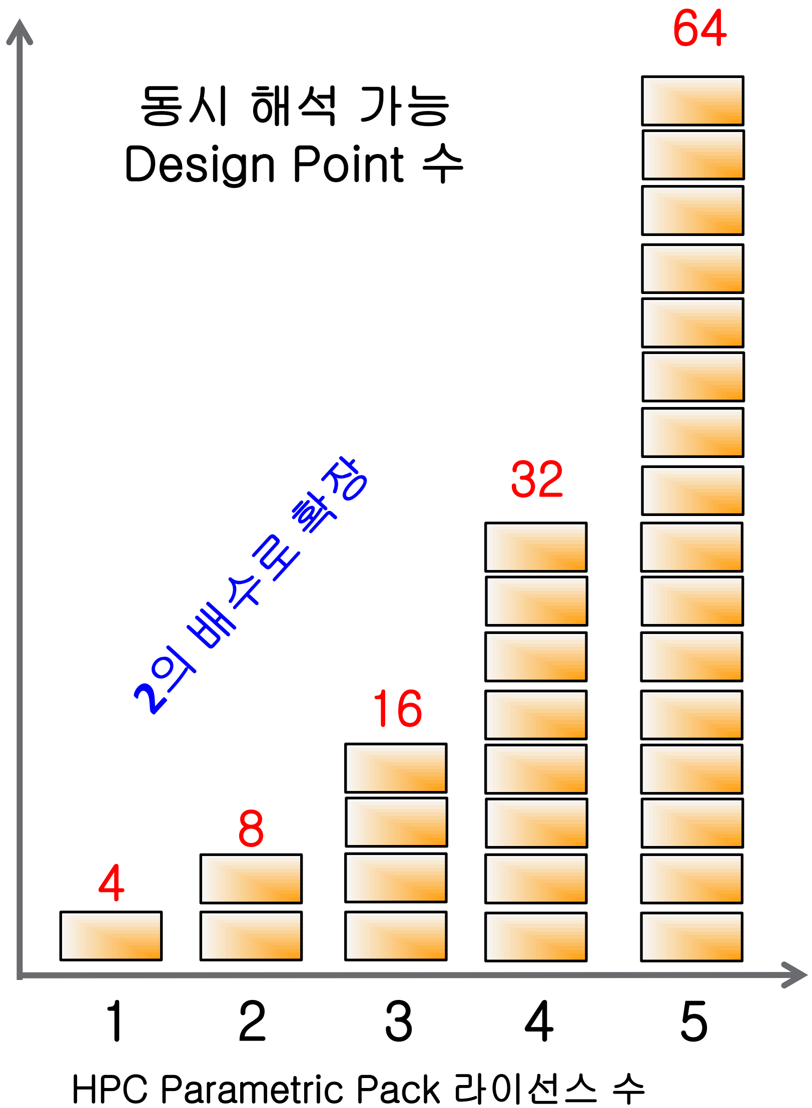

- 하나의 **ANSYS HPC Parametric Pack** 라이선스는 추가적인 응용 프로그램 라이선스가 없어도 **동시에 4개**의 **Design Point**를 해석 가능

- **HPC Parametric Pack** 라이선스가 추가될 때 마다 동시 계산 가능한 **Design Point**의 수를 **2배**로 확장

ANSYS HPC Parametric Pack
예 : Semi-submersible의 구조해석

해석 문제

- 유체정수압과 중력이 가해지는 **Beam과 Shell** 요소를 사용하는 **Semi-submersible**의 정하중해석
- 설계 목적 : 총 무게와 유효 응력을 최소화
- 입력 변수 : **Pontoon** 두께**, Base column** 두께
  (16 design points)
- 상세 사항 :
  - **232,583** 절점**, 230,770** 요소
  - 하드웨어 **: Dell workstation with dual Intel® Xeon® E5-2690 (2.90 GHz, 16 cores), 256 GB memory, all jobs running 2 cores**

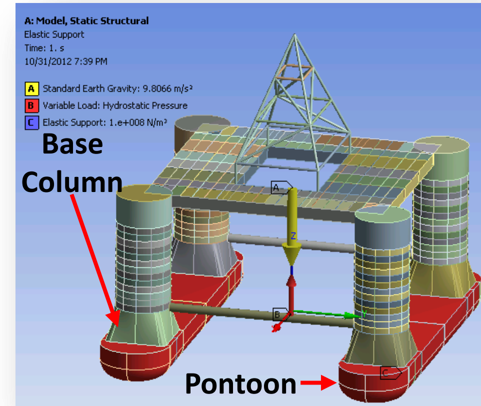

해석에 사용된 라이선스

- **1 ANSYS Mechanical**
- **2 ANSYS HPC Parametric Packs**

X 8

결과/이점

- 순차적인 해석 보다 약 **6**배 속도 향상
- 쉽고 자동화된 작업

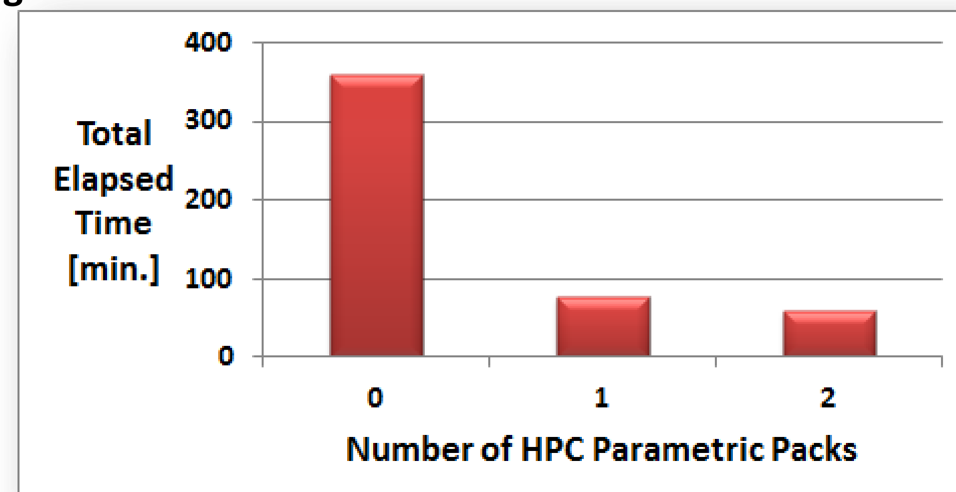

Acknowledgment: Paul Schofield and Jiaping Zhang, ANSYS Houston

ANSYS HPC Parametric Pack
예 : Mixing Vessel의 유동해석

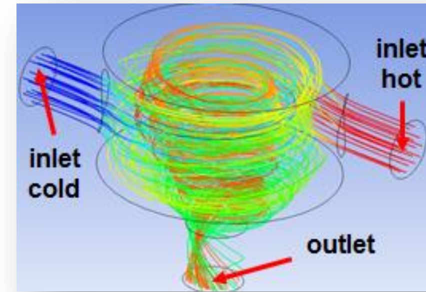

해석 문제
- **Mixing vessel 유동해석**
- **설계 목적 :**
  - **출구에서 온도 분포와 Vessel에서 압력 강하를 최소화되기 위해 운전 조건 안에서 입구 속도를 최적화**
- **입력 변수 : Inlet(cold, hot)에서의 유체 속도 (8 Design Points)**
- 상세 사항 :
  - **K-Epsilon Model with Standard Wall Functions**
  - **52,000 절점, 280,000 요소**
  - 하드웨어 : **HP workstation with dual Intel® Xeon® E5-2687W (3.10 GHz, 16 cores), 128 GB memory**

Outline of All Parameters

	A	B	C	D
1	ID	Parameter Name	Value	Unit
2	Input Parameters			
3	Fluid Flow (FLUENT) (A1)			
4	P1	velocity-1	1	m s^-1
5	P2	velocity-2	2.5	m s^-1
*	New input parameter	New name	New expression	
7	Output Parameters			
8	Fluid Flow (FLUENT) (A1)			
9	P3	pressuredrop	15.83	Pa
10	P4	tempspread	0.31104	K
*	New output parameter		New expression	
12	Charts			

해석에 사용된 라이선스
- **1 ANSYS Fluent**
- **2 ANSYS HPC Parametric Packs**

X 8

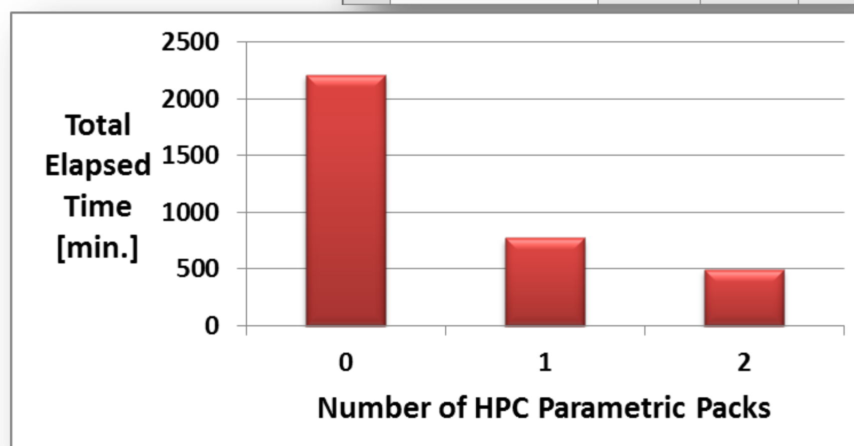

결과/이점
- 순차적인 해석 보다 약 **4.5**배 속도 향상
- 쉽고 자동화된 작업

Acknowledgment: Paul Schofield and Jiaping Zhang, ANSYS Houston

GPU 사용을 위한 라이선스
하나의 HPC 가 하나의 GPU를 위해 요구

라이선스 예 :

유효한 구성의 예 :

1 x ANSYS HPC Pack
총 8 HPC 작업 (최대 4 GPU)

6 CPU Cores + 2 GPUs

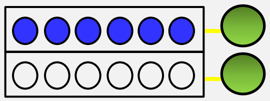

4 CPU Cores + 4 GPUs

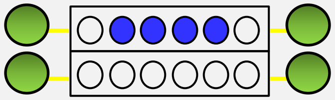

2 x ANSYS HPC Pack
총 32 HPC 작업 (최대 16 GPU)

.
.
.
.
.

24 CPU Cores + 8 GPUs

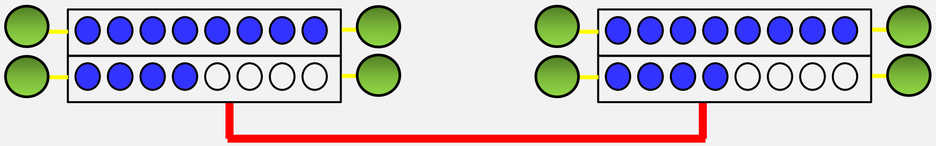

(Total Use of 2 Compute Nodes)

(모든 형태에 적용: HPC, HPC Pack, HPC Workgroup, HPC Enterprise 등)

지원되는 GPU

제조사	지원 하드웨어	지원 운영체제
nVIDIA	Tesla 시리즈 - Fermi C2050, C2070 - Kepler K20, K40, K60 Quadro 6000 Quadro K5000 Quadro K6000	MS Windows 64bit 계열 Linux 64bit 계열
Intel	Xeon Phi 7120 Xeon Phi 5110 Xeon Phi 3120	Linux 64bit 계열 (Red Hat Linux 5.x 제외)

ANSYS Mechanical의 HPC와 GPU 성능

TAESUNG S&E, INC.

ANSYS Mechanical의 HPC와 GPU 성능

Part.1 ANSYS Mechanical에서 GPU 가속

GPU 가속을 위한 솔버
Mechanical

▪ 지원되는 솔버

해석 종류	솔버
Static	Sparse, PCG, JCG
Buckling	Block Lanczos, Subspace
Modal	Block Lanczos, Subspace, PCG Lanczos, QR Damped, Unsymmetric, Damped
Harmonic	Full method(Sparse)
Transient	Full method(Sparse, PCG, JCG)

▪ 솔버별 **GPU** 사용 제한

병렬처리	솔버 종류	**GPU 사용 제한**
SMP (Shared memory parallel)	Sparse	단지 1개의 GPU만 활성화
	PCG, JCG	모든 요청된 GPU 활성화
DMP (Distributed memory parallel)	Sparse	GPU 사용 수가 CPU 사용수를 초과할 수 없음
	PCG, JCG	모든 요청된 GPU 활성화

▪ GPU 가속이 지원되지 않는 솔버 사용시 GPU가속을 활성화하지 않고 해석 진행
▪ Intel Xeon Phi는 상단의 푸른색 글씨 솔버만 지원(SMP 솔버만 지원)

GPU 가속 활성화
Mechanical

- **Mechanical APDL**

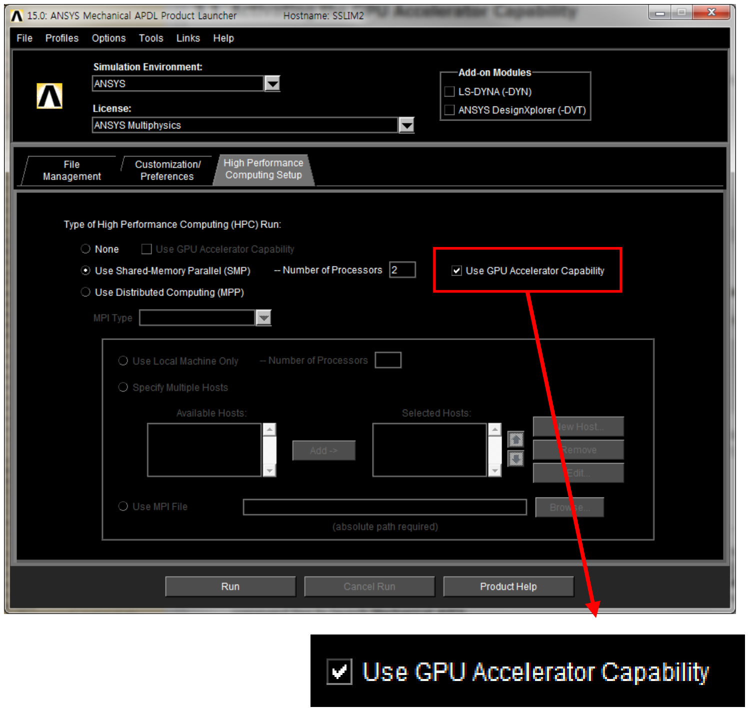

OR

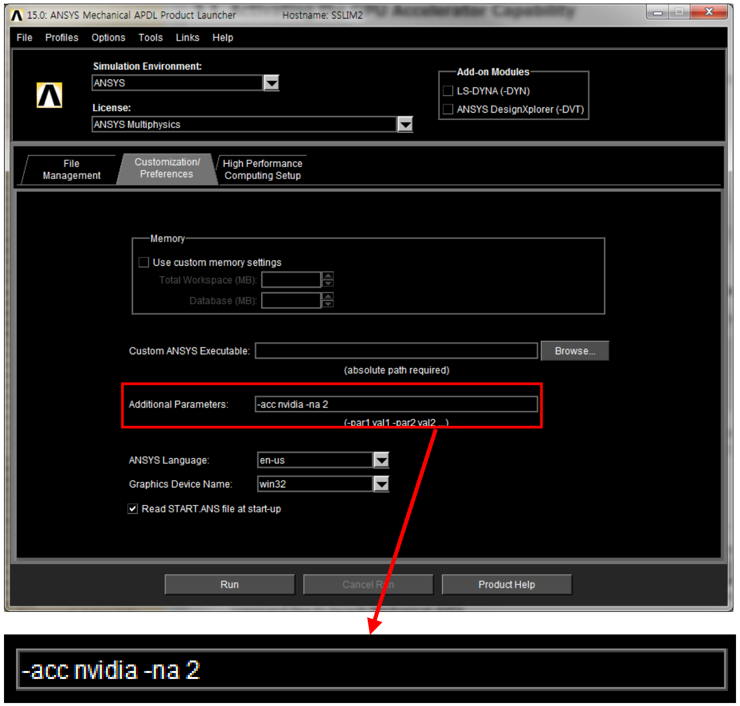

ansys150 –acc nvidia –na N
혹은
ansys150 –acc intel –na N

- GPU 사용시 -na 옵션을 지정하지 않은 경우 Default로 –na 1이 지정
- Input 파일 내부에 ACCOPTION 명령어 사용으로 GPU 활성화 가능

GPU 가속 활성화
Mechanical

- **Workbench Mechanical**

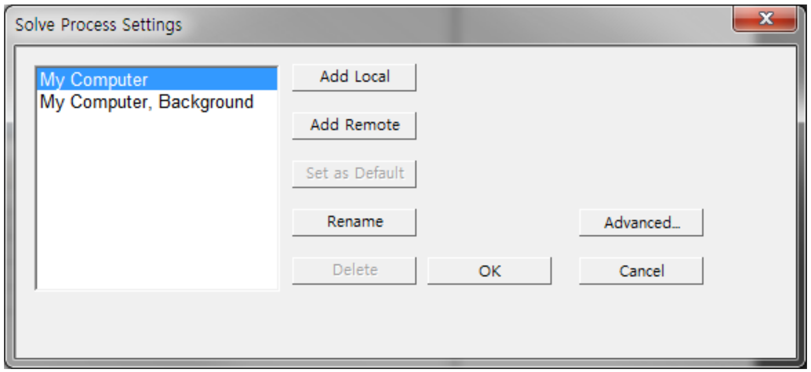

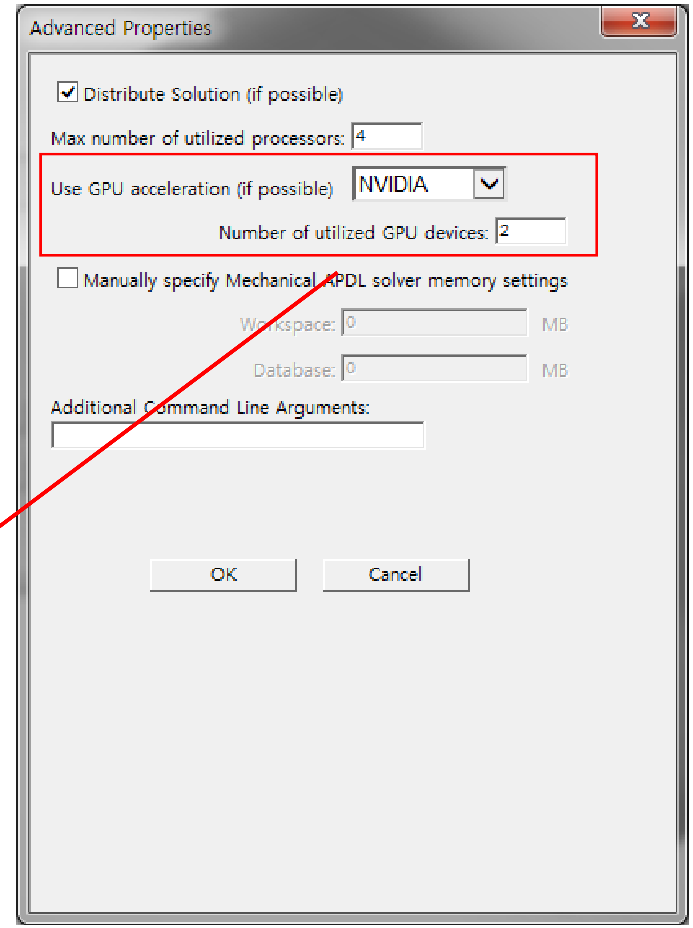

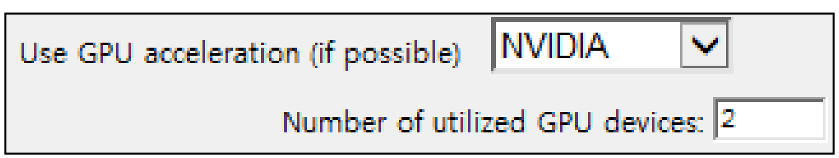

ANSYS Mechanical의 HPC와 GPU 성능

Part.2 ANSYS Mechanical 벤치마크

ANSYS Mechanical에서 벤치마크 테스트
하드웨어 구성

부품	사양
CPU	Intel Xeon E5-2670 2.6GHz - 2개
메모리	128G (1,600MHz)
저장장치	SAMSUNG SSD 840 Pro 256GB - 4개, RAID0
	7200rpm, 1TB
GPU	nVIDIA Tesla K40C – 4개

ANSYS Mechanical에서 벤치마크 테스트

테스트 해석 파일

SP-1	
Model Name	Universal Joint
Analysis Type	Static, Nonlinear, Structural
Solver Choice	Sparse, Real-Value, Symmetric
Element Type	SOLID186 (Nodes : 144,218, Elements : 99,609)
FE Model	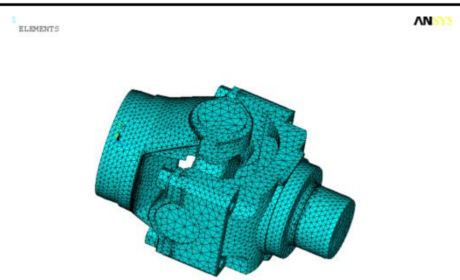 17 Iterations Small sized job 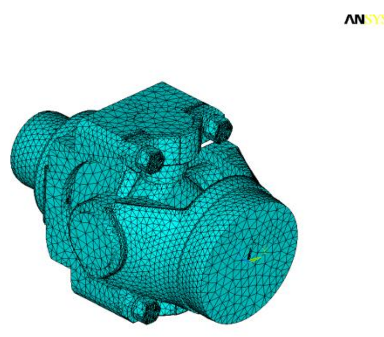

SP-2	
Model Name	Peltier Cooling Block
Analysis Type	Static, Nonlinear, Thermal-Electric Coupled-Field
Solver Choice	Sparse, Real-Value, Symmetric
Element Type	SOLID226, SOLID,227 (Nodes : 277,715, Elements : 152,717)
FE Model	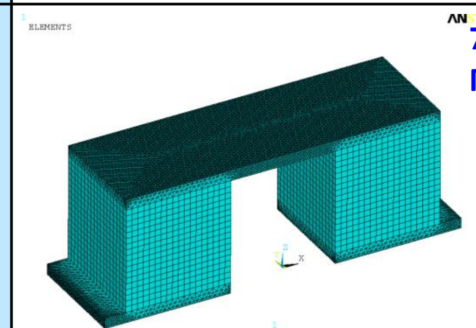 7 Iterations Medium sized job 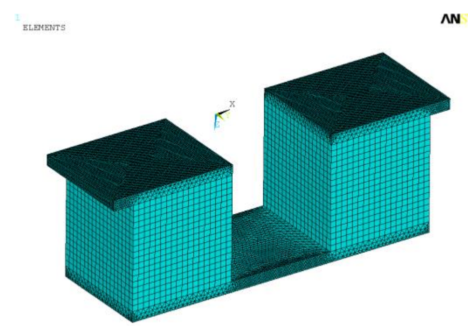

ANSYS Mechanical에서 벤치마크 테스트
테스트 해석 파일

SP-3	
Model Name	Semi-Submersible
Analysis Type	Transient, Nonlinear, Structural
Solver Choice	Sparse, Real-Value, Symmetric
Element Type	SHELL281, BEAM189 (Nodes : 396,795, Elements : 136,867)
FE Model	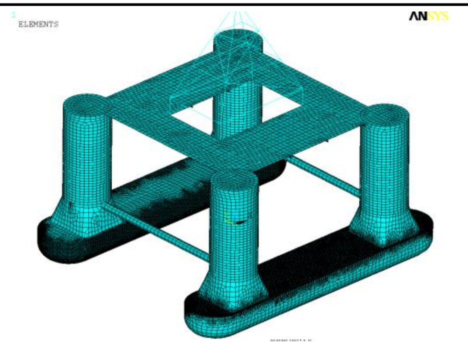 11 Iterations Large sized job 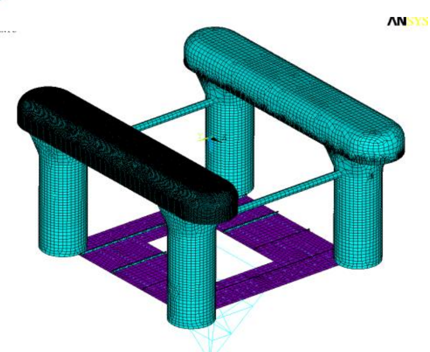

SP-4	
Model Name	Carrier
Analysis Type	Harmonic, Linear, Structural
Solver Choice	Sparse, Complex-Value, Symmetric
Element Type	SOLID187 (Nodes : 332,362, Elements : 222,005)
FE Model	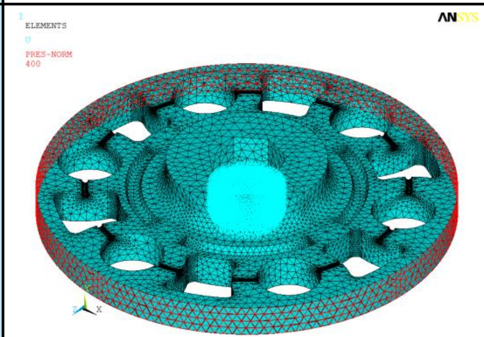 4 Frequencies Large sized job 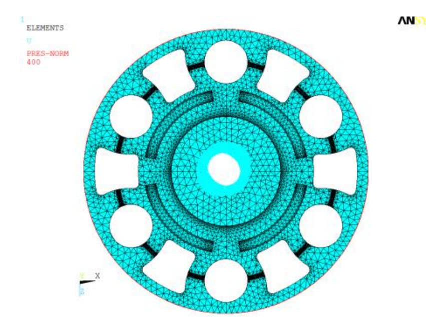

ANSYS Mechanical에서 벤치마크 테스트
테스트 해석 파일

SP-5	
Model Name	Turbine
Analysis Type	Static, Nonlinear, Structural
Solver Choice	Sparse, Real-Value, Symmetric
Element Type	SOLID187, CONTA174. TARGE170 (Nodes : 715,008, Elements : 483,631)
FE Model	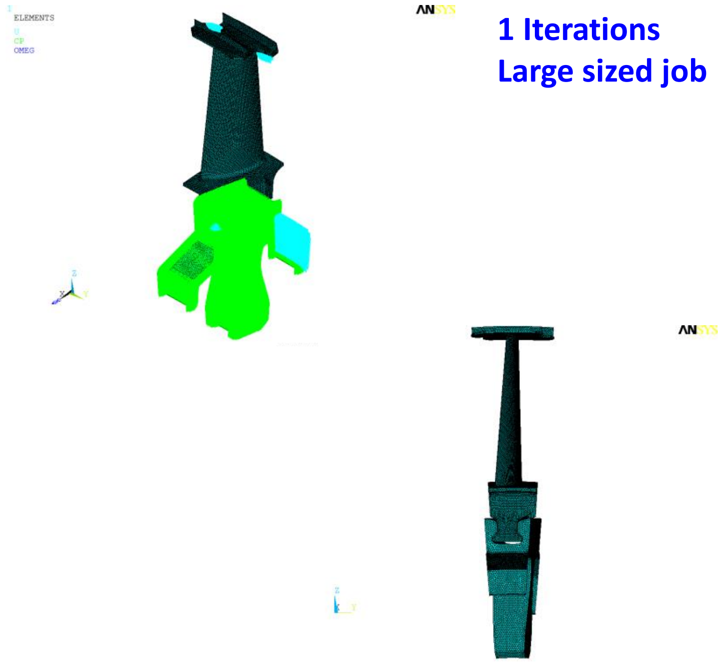 1 Iterations Large sized job

26

ANSYS Mechanical에서 벤치마크 테스트
Block Lanczos vs. Subspace eigen solver

- R15에서 새로운
  Subspace eigen solver 지원
- Subspace eigen solver는
  SMP와 DMP 기술 지원

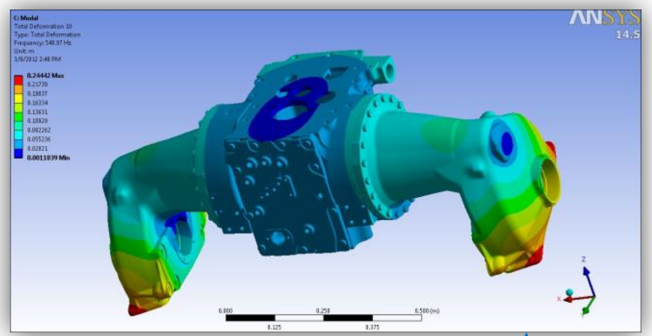

2.09 MDOFs
first 20 modes

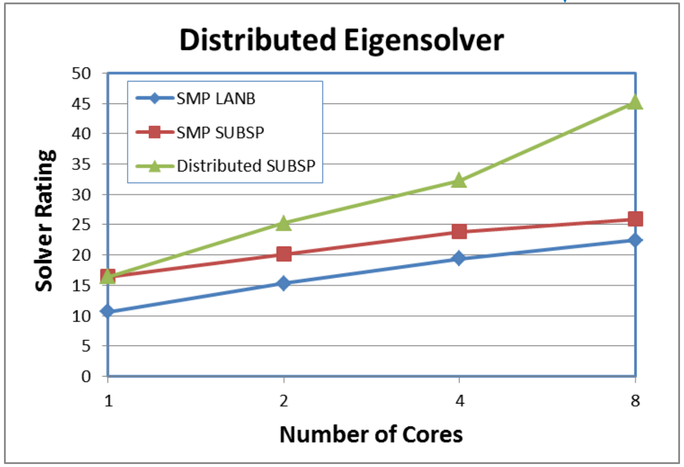

ANSYS Mechanical에서 벤치마크 테스트

Block Lanczos vs. Subspace eigen solver

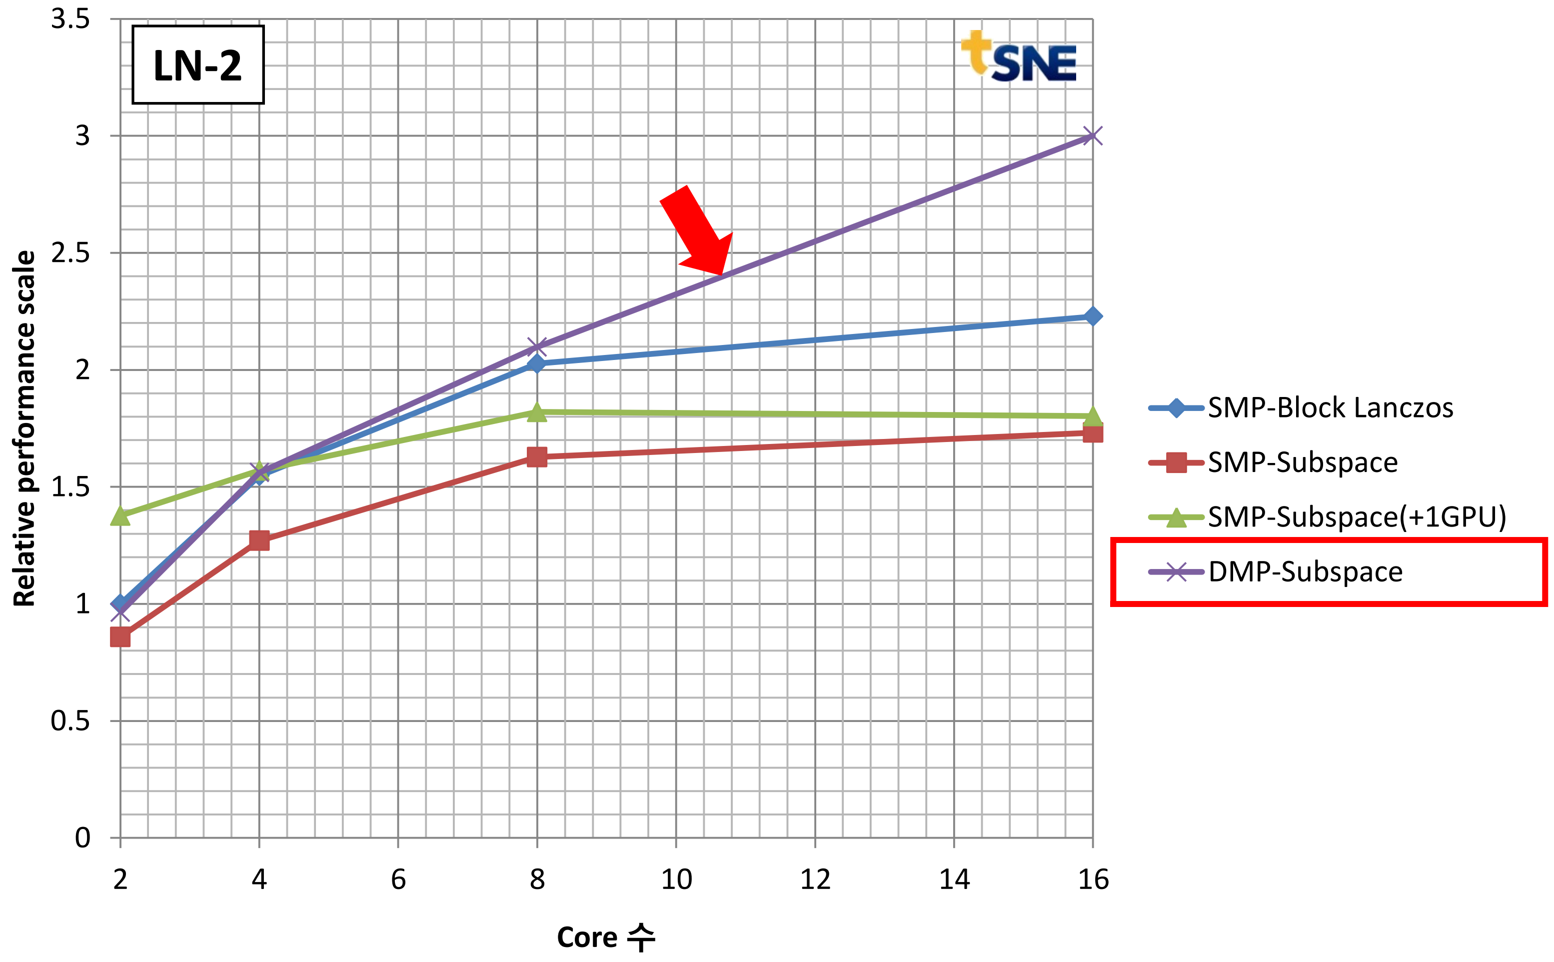

- Subspace eigen solver는 DMP 지원
- SMP 방식은 Block Lanczos가 더 빠름
- DMP 방식의 Subspace eigen solver는 Core 수에 비례하여 속도 향상

ANSYS Mechanical에서 벤치마크 테스트

R14.5 vs. R15.0 Sparse solver

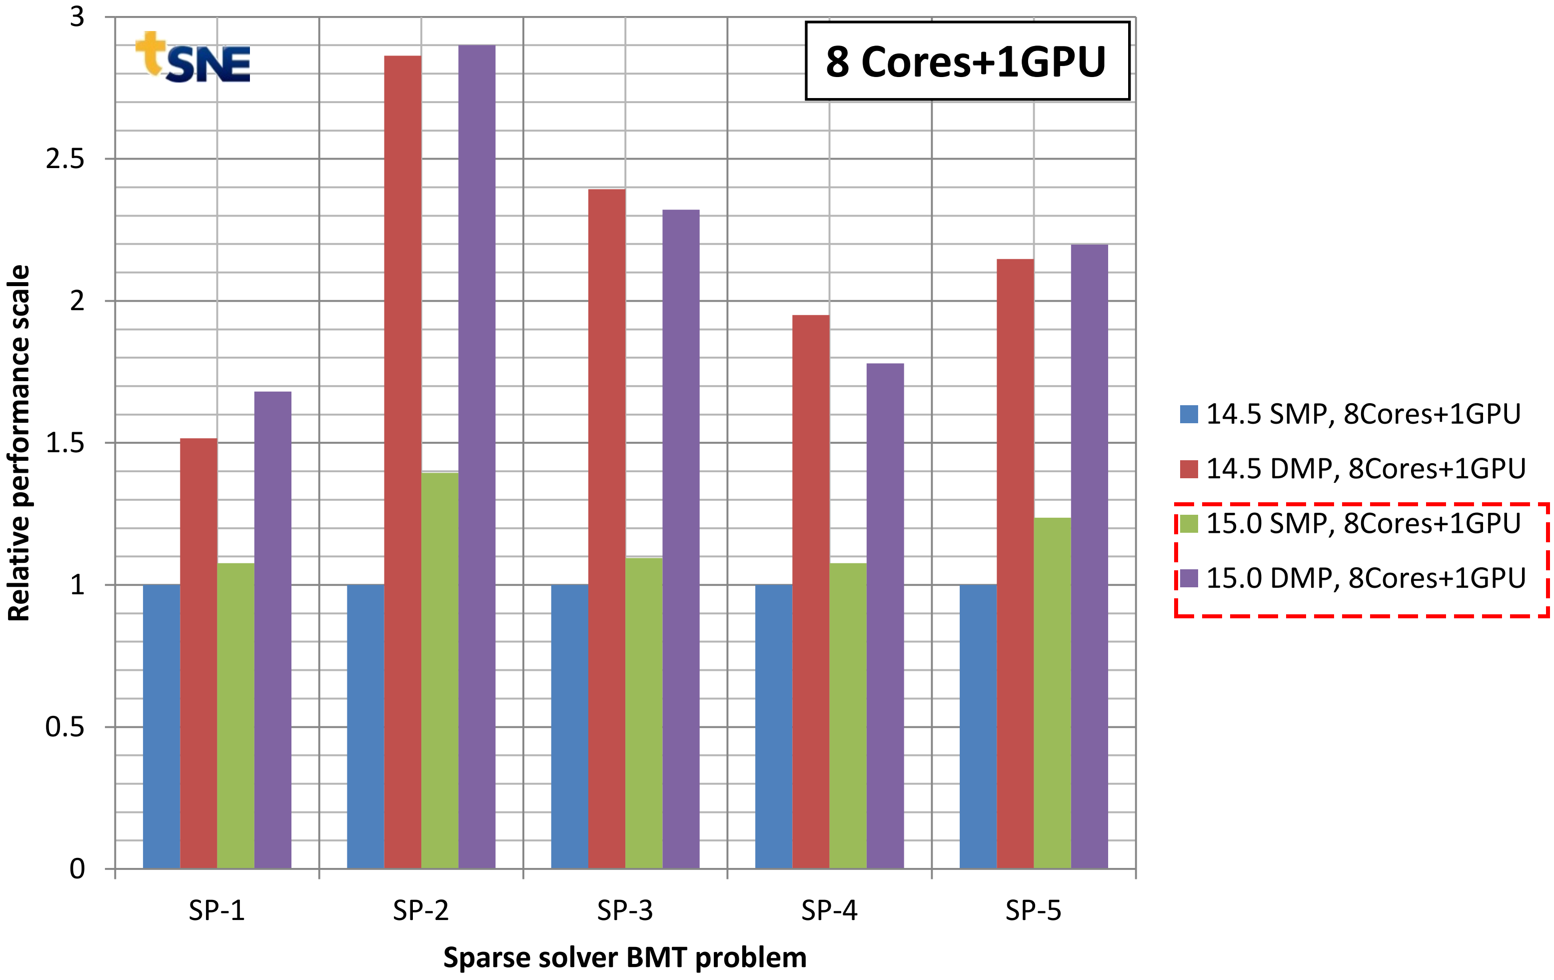

- R14.5에 비해 R15.0 속도 향상
- SMP 보다 DMP가 속도가 빠르고 Core 수가 올라갈 때 유리

ANSYS Mechanical에서 벤치마크 테스트
Shared memory parallel solver와 GPU 가속

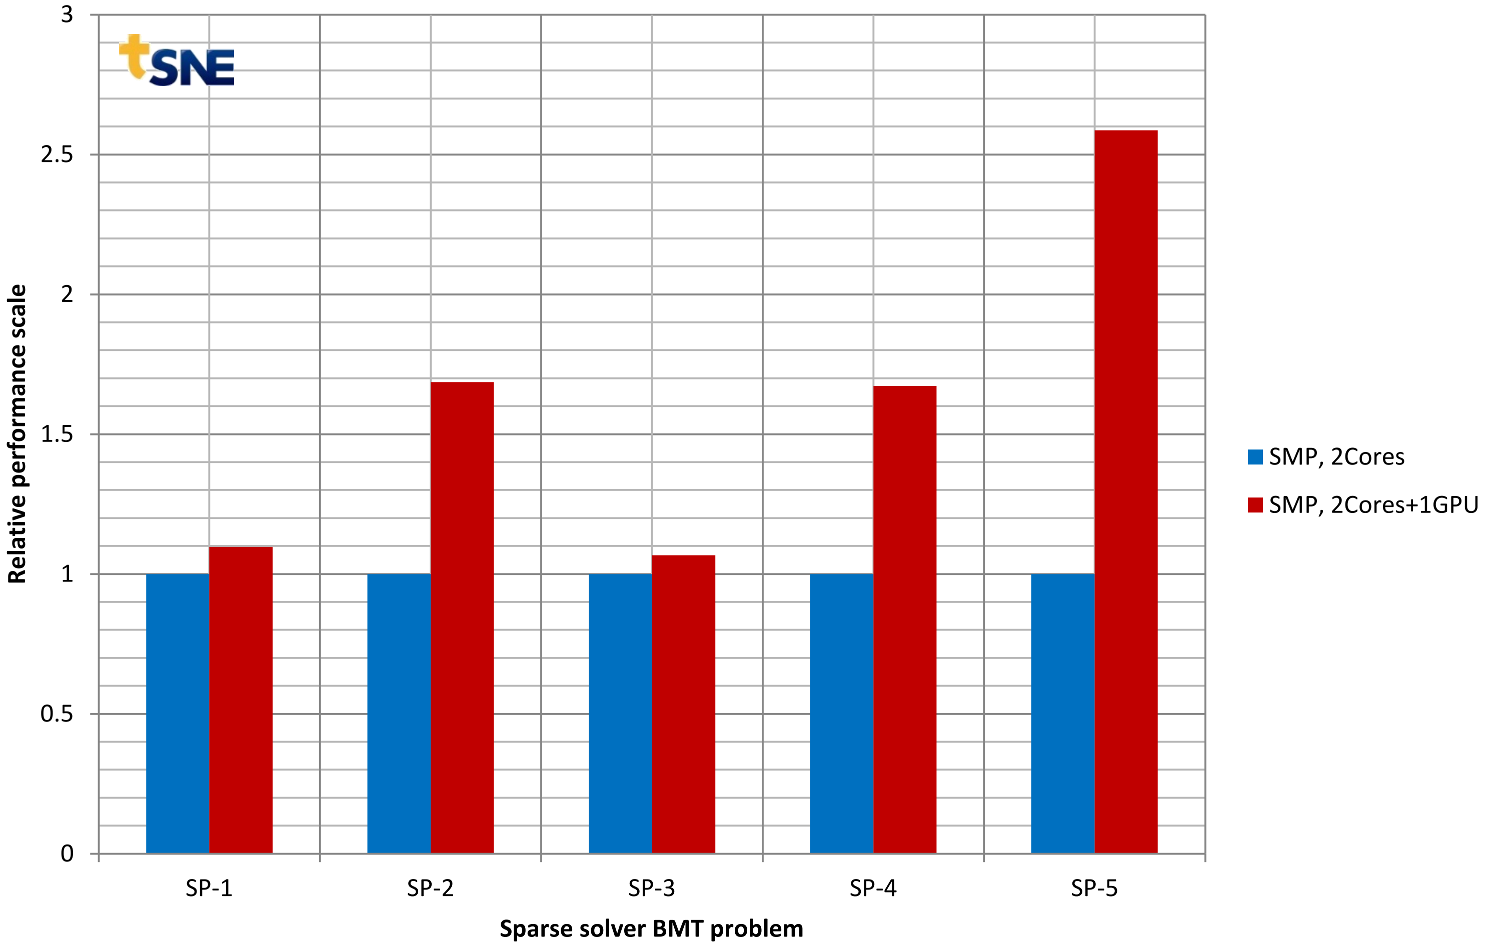

- SMP 솔버에 GPU 가속을 추가할 경우 계산 속도가 향상
- SP-5 문제 형태에 대해서는 단지 하나의 HPC 라이선스 추가로 2.5배 이상의 계산 속도 향상

ANSYS Mechanical에서 벤치마크 테스트
Distributed memory parallel solver와 GPU 가속

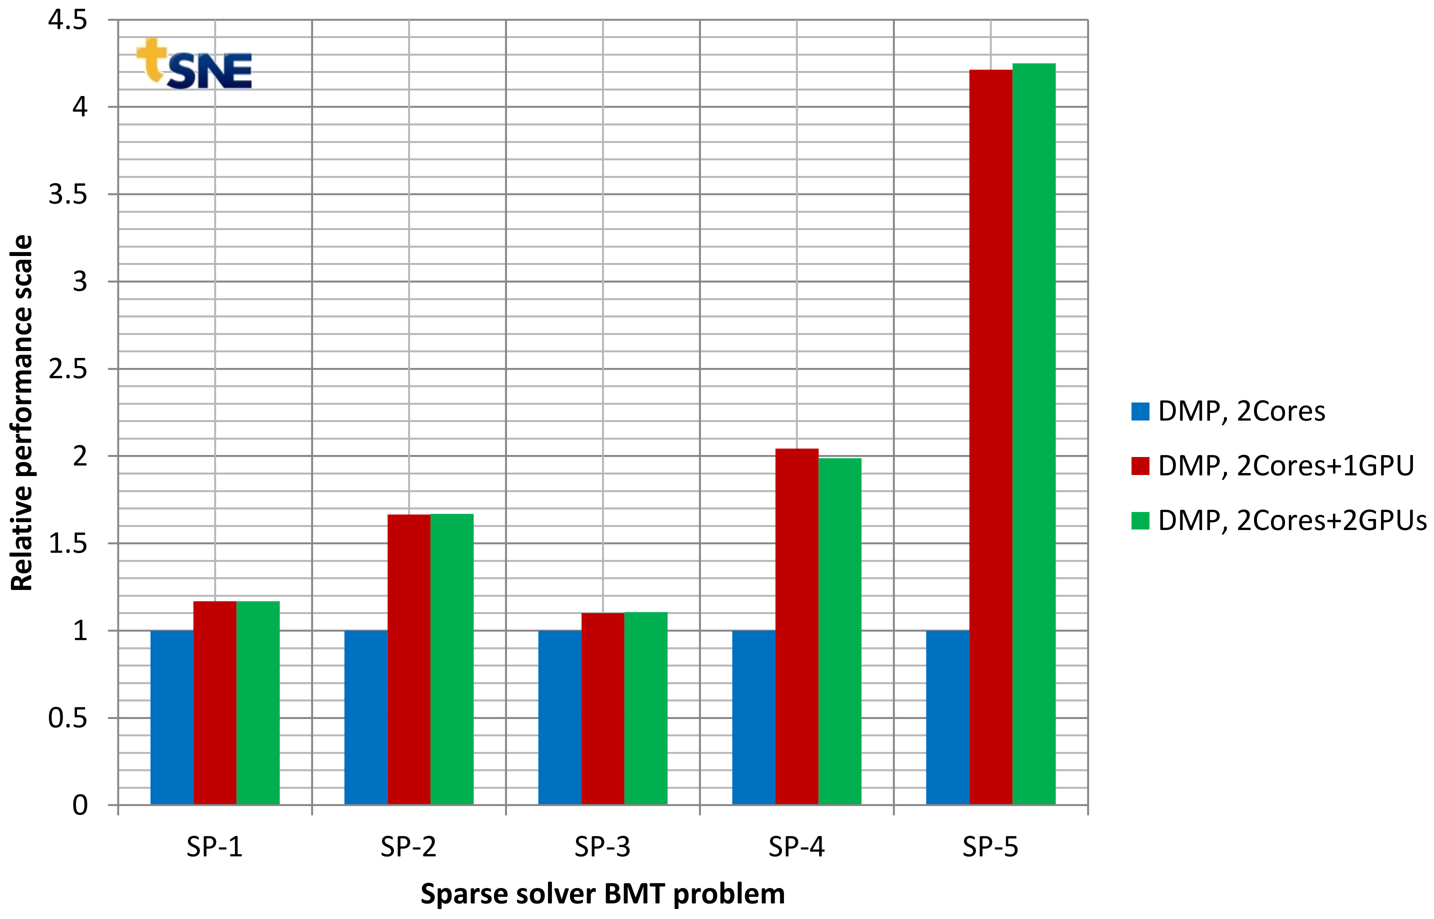

- DMP 솔버에 GPU 가속을 추가할 경우 계산 속도가 향상
- SP-5 문제 형태에 대해서는 단지 하나의 HPC 라이선스 추가로 4배 이상의 계산 속도 향상

ANSYS Mechanical에서 벤치마크 테스트
HPC Pack 라이선스 사용자를 위한 GPU

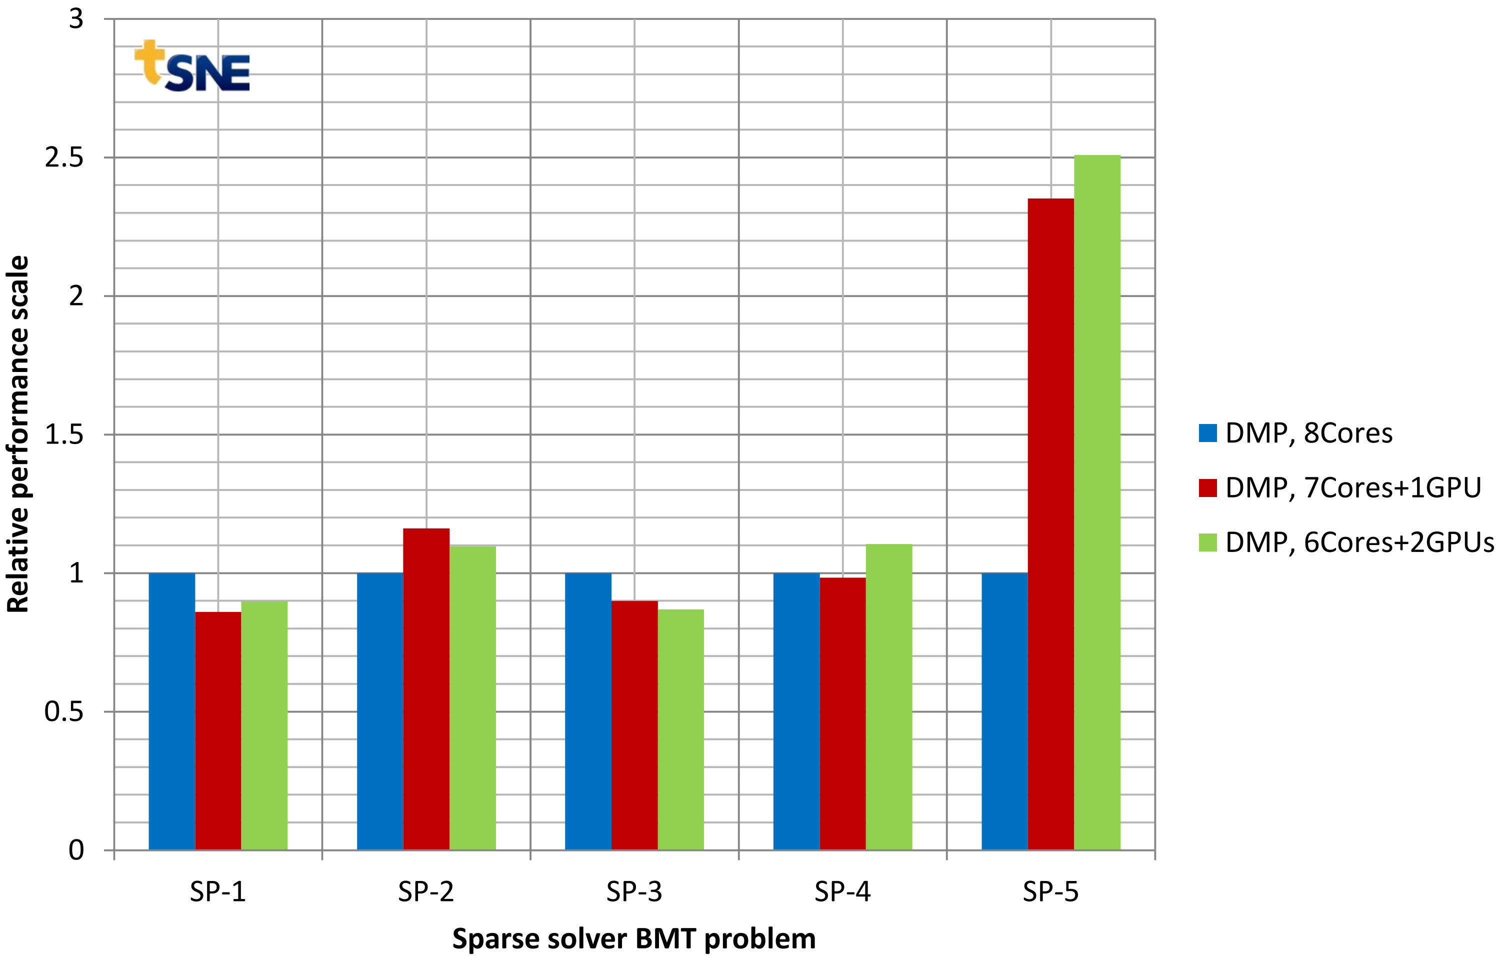

- DMP 솔버 기준 총 8 HPC 작업
- SP-5 문제 형태에 대해서는 GPU를 추가할 수록 성능 향상

ANSYS Mechanical에서 벤치마크 테스트
GPU 가속

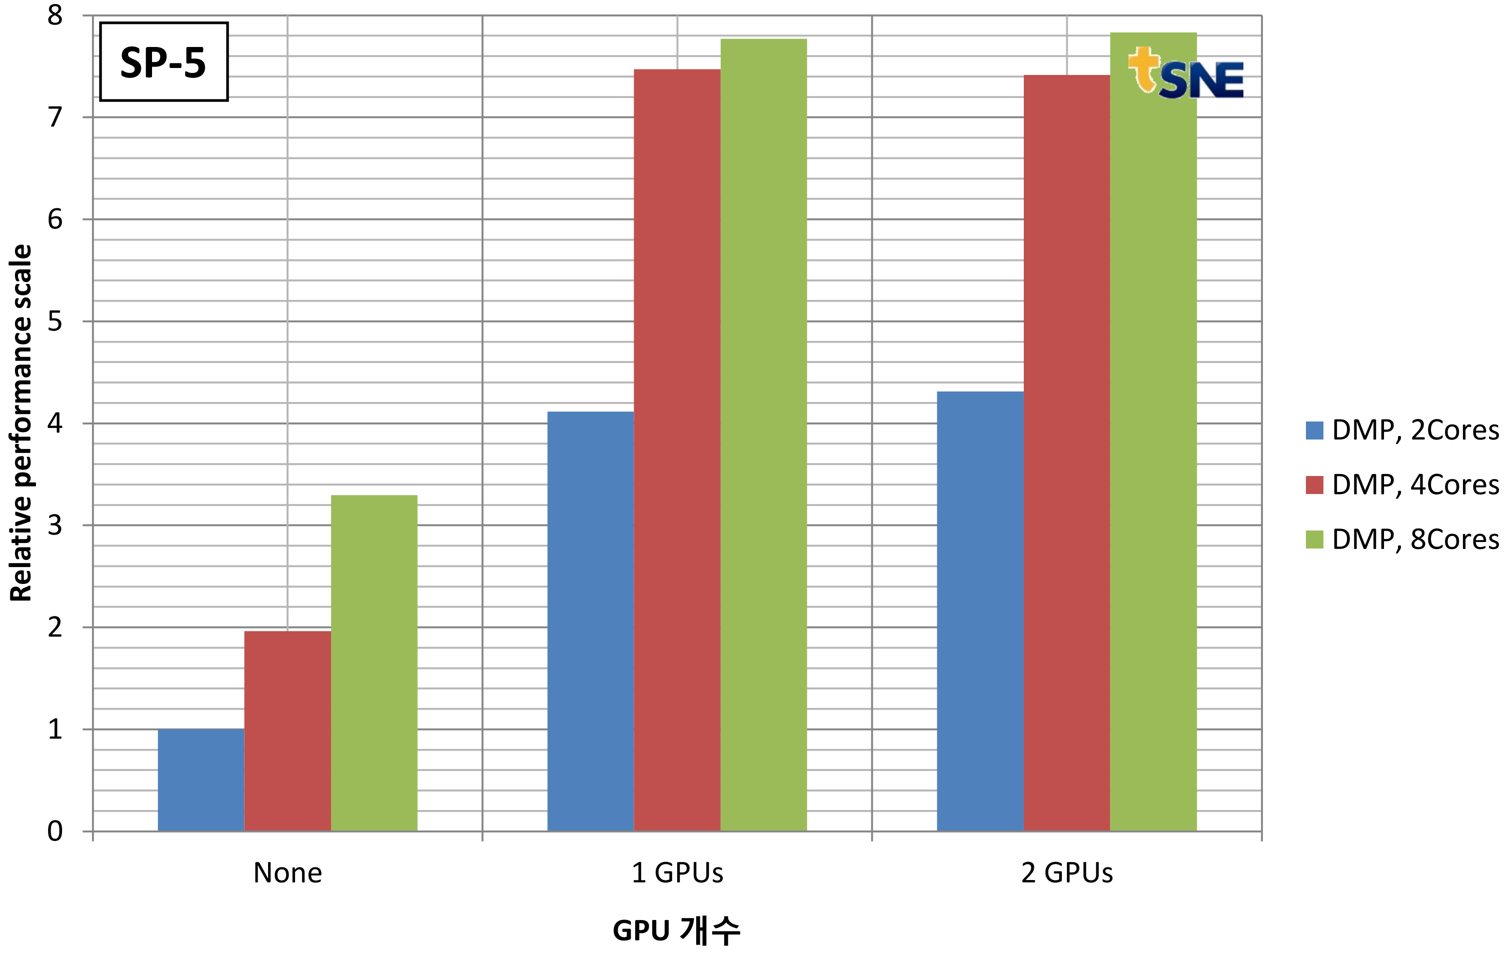

- SP-5 벤치마크 테스트 문제
- GPU가 1개 추가 될 때 속도 향상이 뚜렸함

ANSYS Mechanical에서 벤치마크 테스트

Intel Xeon Phi

Xeon Phi card와 함께 수행된 속도 향상

- **Linux 환경에서 Shared Memory Sparse Solver**

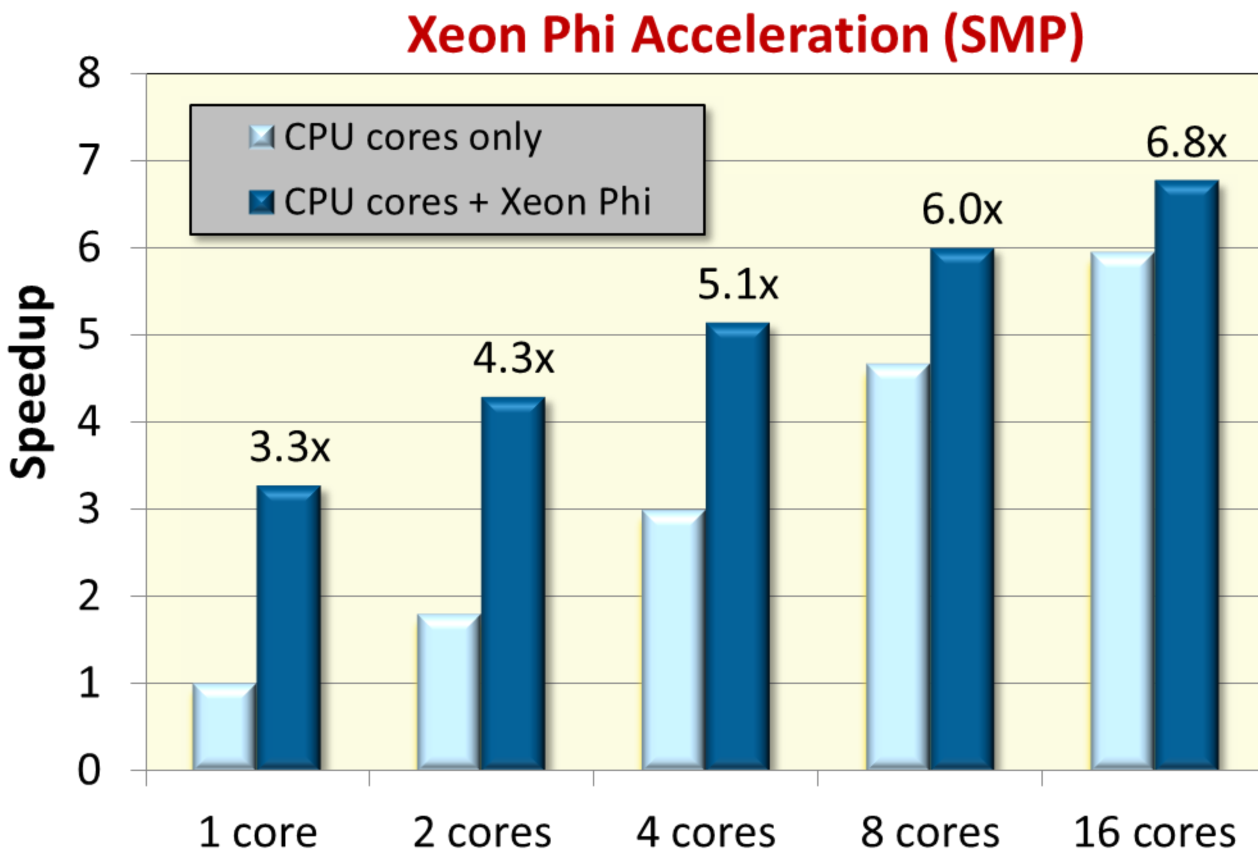

V14sp-5 Model

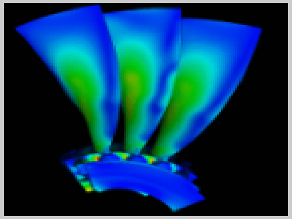

- Turbine geometry
- 2.1 million DOF
- SOLID187 elements
- Static, nonlinear analysis
- One iteration
- Sparse direct solver

Linux workstation (16 Intel Xeon E5-2670 cores @ 2.6 GHz, 1 7120A Xeon Phi, 64 GB RAM)

ANSYS CFD의 HPC와 GPU 성능

ANSYS CFD의 HPC와 GPU 성능

Part. 1 ANSYS CFD R15.0

ANSYS CFD R15.0

Parallel Enhancement

ANSYS CFX : Parallel Scalability

R&D project to improve HPC

- Investigation of various solver parallel scalability limitations
- Industrial benchmarks
  - Single and multi-domain (incl. two-stage radial compressor and six-stage axial compressor)
  - Steady and transient
- Implemented improvements accessible via expert parameter
  - Default setting does not incorporate changes
    - https://sites.google.com/a/ansys.com/cfx-users/cfx-beta-features/release-15/hpc/parallel-optimization

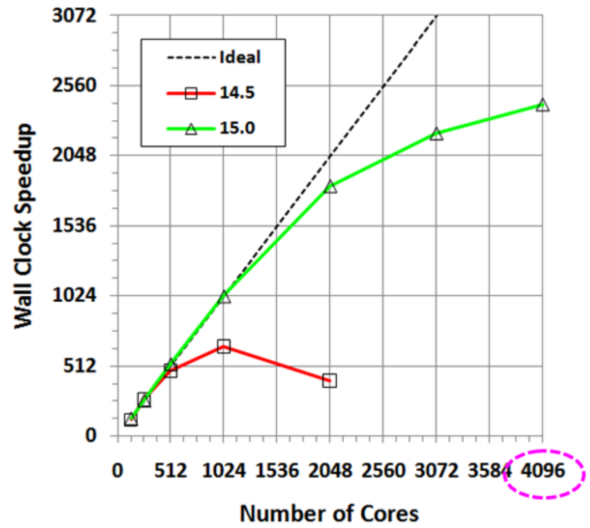

69 seconds, (89% efficient @2048 cores)

~4X faster than V14.5

311 seconds (20% efficient @2048 cores)

Solver wall clock speed-up on 150M node intake case

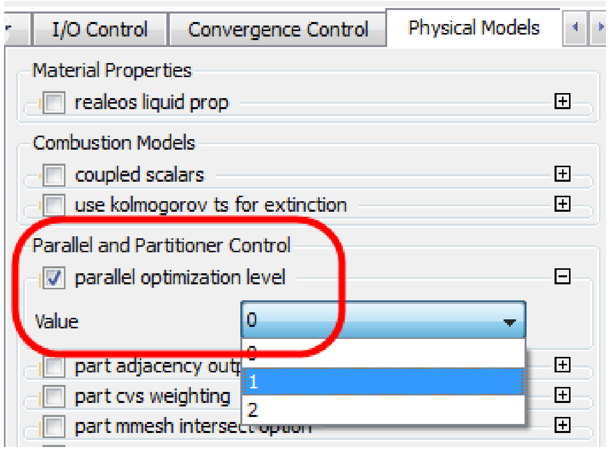

ANSYS CFX : Parallel Scalability

Industrial benchmark application

- 6-Stage Axial Compressor
- 13M nodes, 14 Domains, 12 Mixing Planes

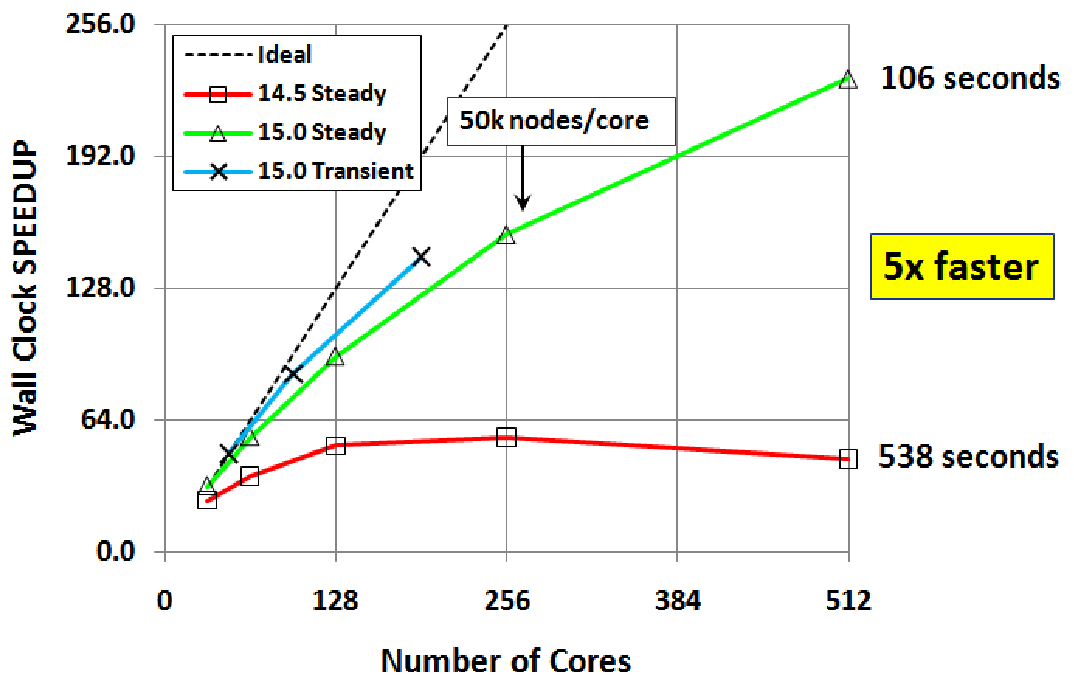

Courtesy Siemens AG, Mülheim, Germany, ASME IGTI Paper GT2013-94639

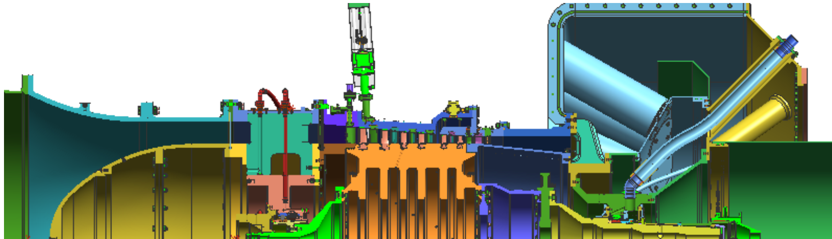

ANSYS FLUENT : Parallel Scalability

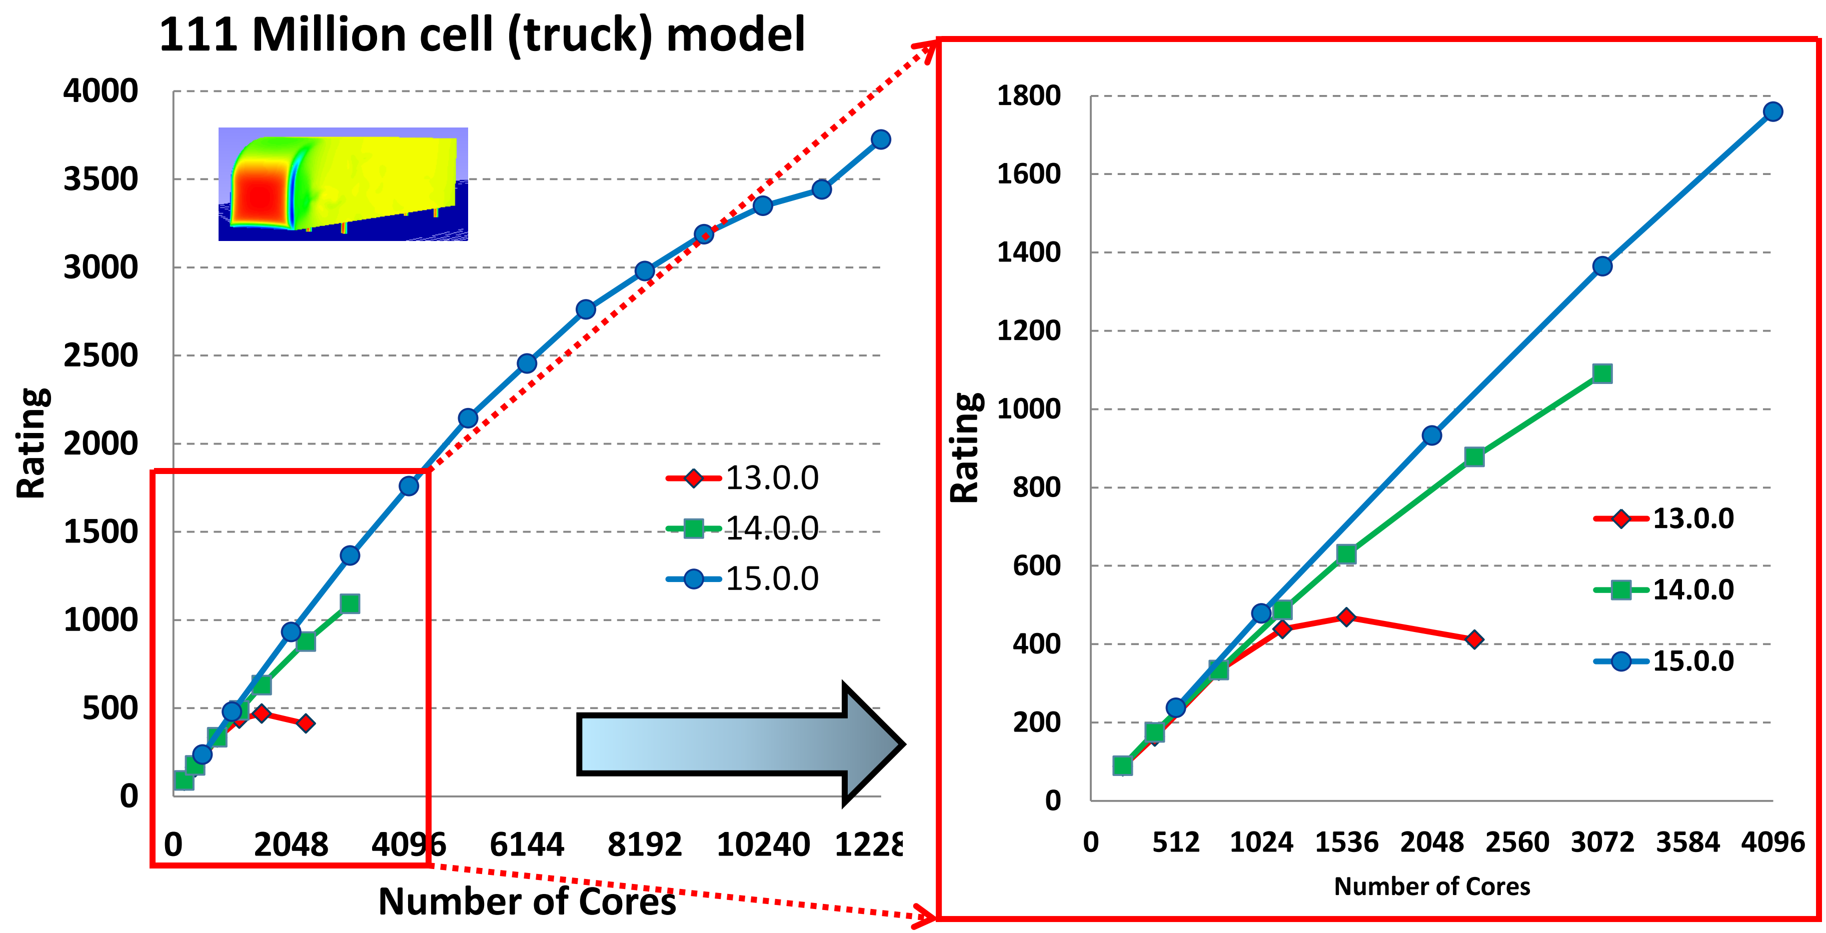

Scalable at 10,000 cells per core!

ANSYS FLUENT : Parallel Scalability

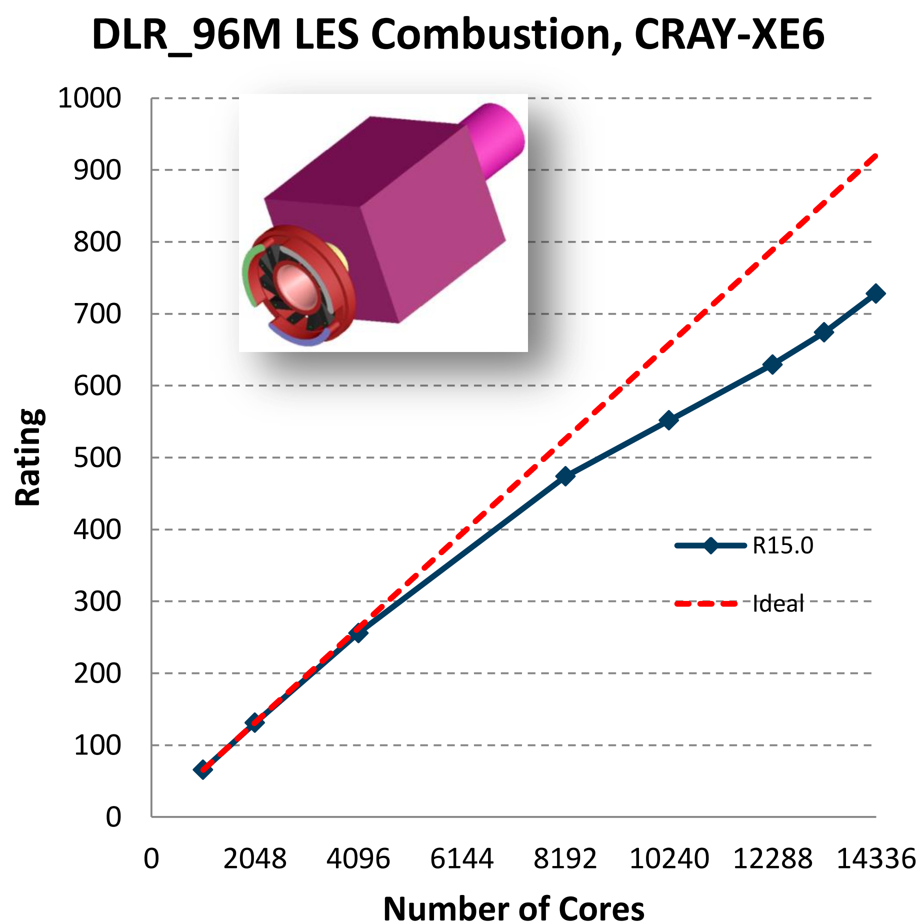

- Gas phase combustion
- Thickened Flame Model(TFM) with Finite Rate Chemistry
- Pressure based coupled
- 96 million cells, hex-core mesh
- ***84%*** **efficiency at 10,240 cores**
  - ➢ **Less than 10,000 cells per core!**

Number of Cores	R15.0	Efficiency
1024	65.7518	100
2048	131.209	99.78
4096	256.072	97.36
8192	473.856	90.08
10240	551.844	83.93

ANSYS FLUENT : Parallel Scalability for Discrete Phase

New 2-Domain method

- Balance the continuous and discrete phase independently
- Over 2x improvement seen for 512-way parallel

Improved scalability for hybrid method

- Default method for parallel particle tracking

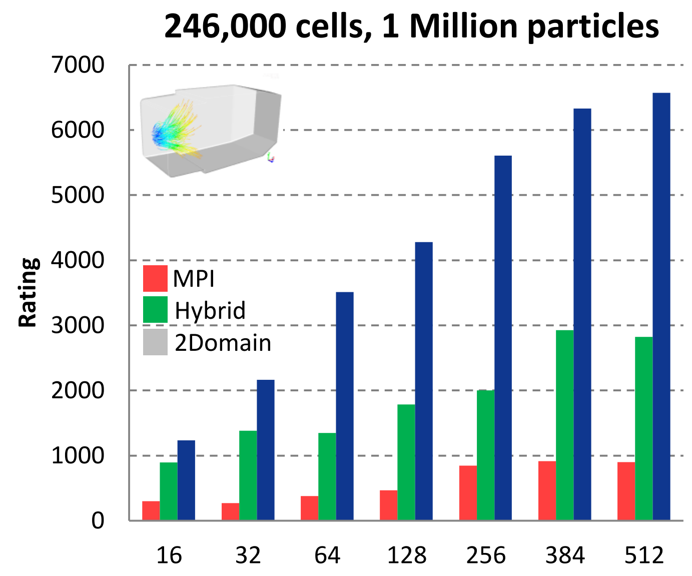

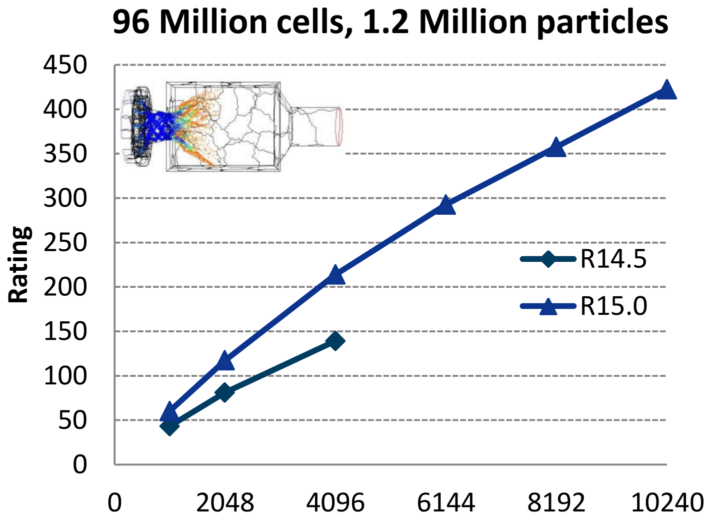

ANSYS FLUENT : Parallel File I/O and Startup

More efficient parallel I/O and startup

- Case read time reduced significantly at high core counts
- Start-up time for 8192-way parallel reduced from 30 minutes to 30 seconds

Effective configuration of parallel processes

- Use different number of processes for meshing and solve modes

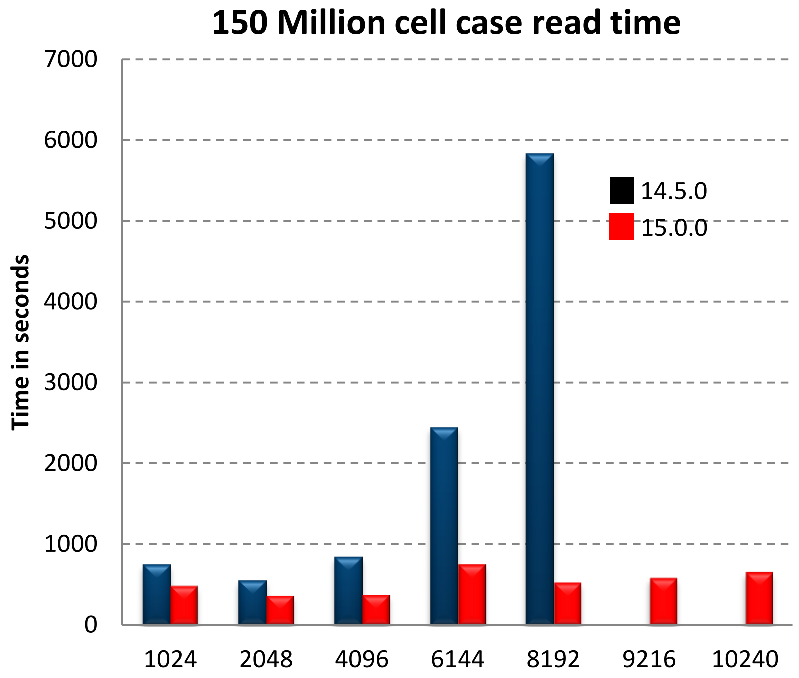

ANSYS FLUENT : Other Parallel Enhancements

Improved parallel error handling

- Ability to restore running simulations to a usable state after a crash

Faster solutions using GPUs

- Accelerated AMG solver performance for 3D coupled pressure-based solver cases

Internal Flow Steady Computation								
		Fluent		NVAMG				
		AMG	Total	AMG		Total		
Serial	SP	352	558	74	4.76x	289	1.93x
	DP	485	777	97	5.00x	417	1.86x
T12	SP	71	100	86	0.83x	114	0.88x
	DP	132	179	119	1.11x	166	1.08x

Internal Flow Unsteady Computation							
		Fluent		NVAMG			
		AMG	Total	AMG		Total	
Serial	SP	1983	3168	394	5.03x	1582	2.00x
	DP	2832	4593	517	5.48x	2263	2.03x
T12	SP	496	653	454	1.09x	611	1.07x
	DP	933	1221	630	1.48x	920	1.33x

Support for Intel Many-Integrated-Core (MIC) (β)

- Intel Xeon Phi

ANSYS CFD의 HPC와 GPU 성능

Part. 2 ANSYS FLUENT with GPU

ANSYS FLUENT R15.0

GPU Parallel

MiruWare System & Solution

NVIDIA

TAESUNG S&E, INC.

How to Enable NVIDIA GPUs in FLUENT 15.0

Windows :

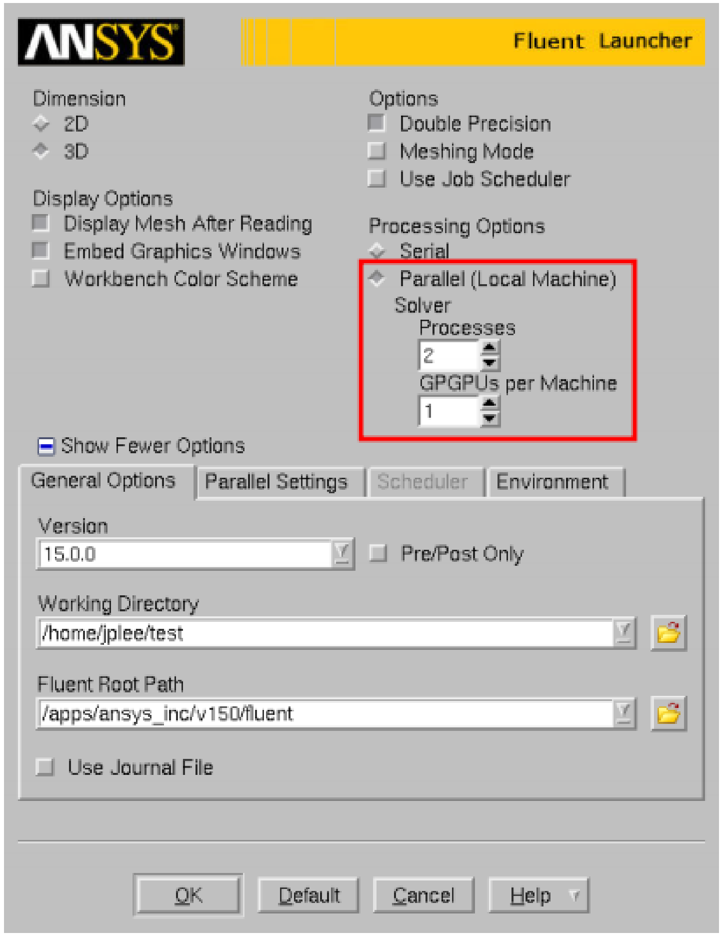

Linux :

Fluent 3ddp -t2 –gpgpu=1

Cluster specification :

nprocs = Total number of fluent process

M = Number of machines

ngpgpus = Number of GPUs per machine

Requirement 1:

nprocs mod *M* = 0

: Same number of solver processes on each machine

Requirement 1:

(nprocs / M) mod *ngpgpus* = 0

: Number of processes should be an interger multiple of GPUs

Cluster Specification Examples

Single-node configurations :

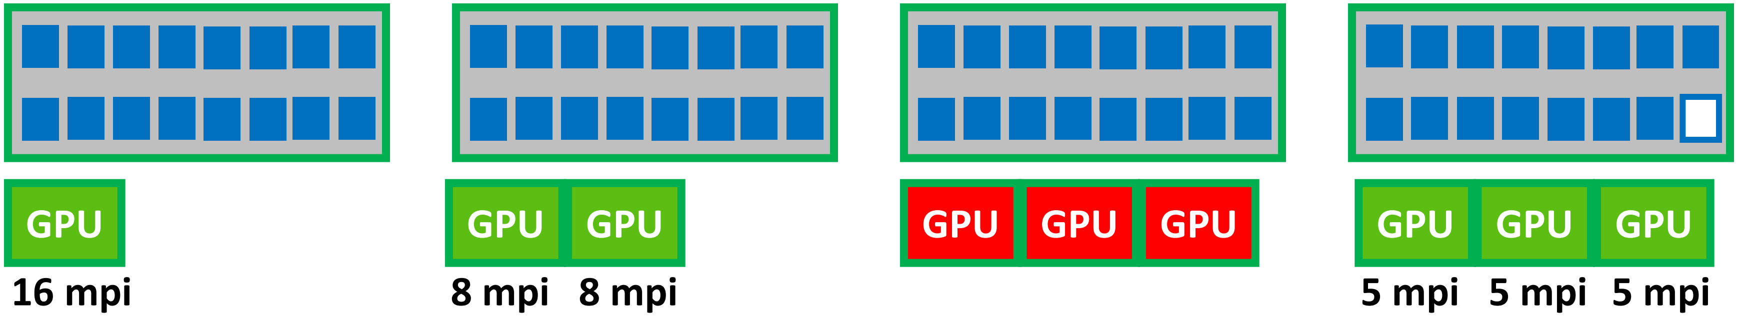

Multi-node configurations :

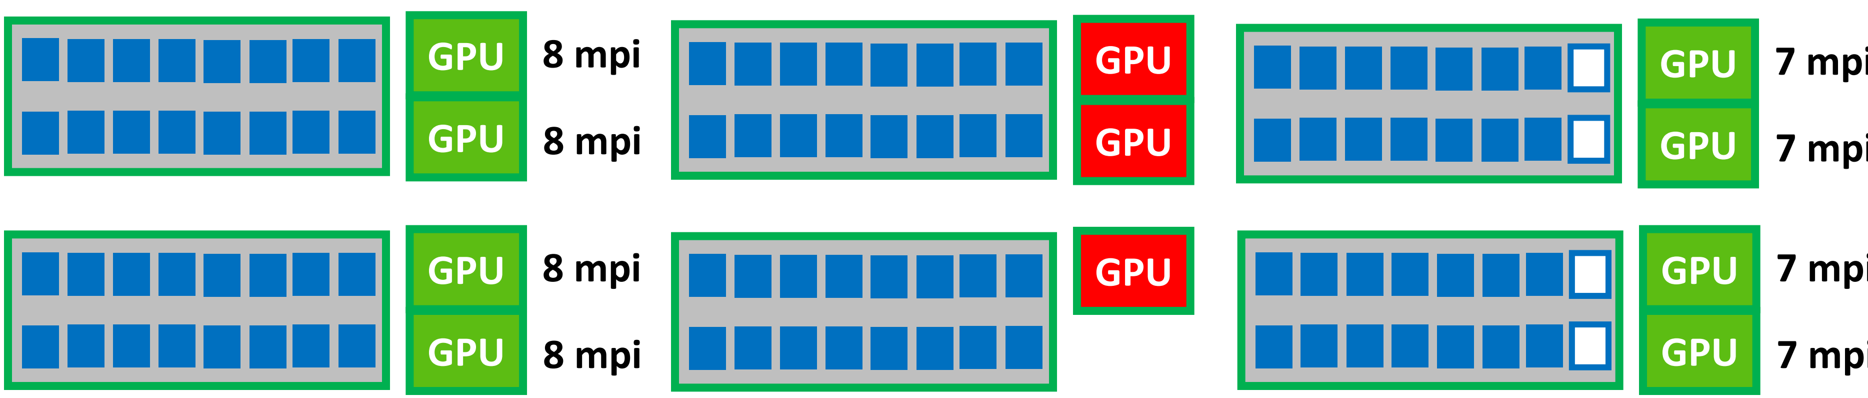

ANSYS Fluent 14.5 GPUs Acceleration

Preview of ANSYS Fluent 14.5 Performance – by ANSYS, Aug 2012

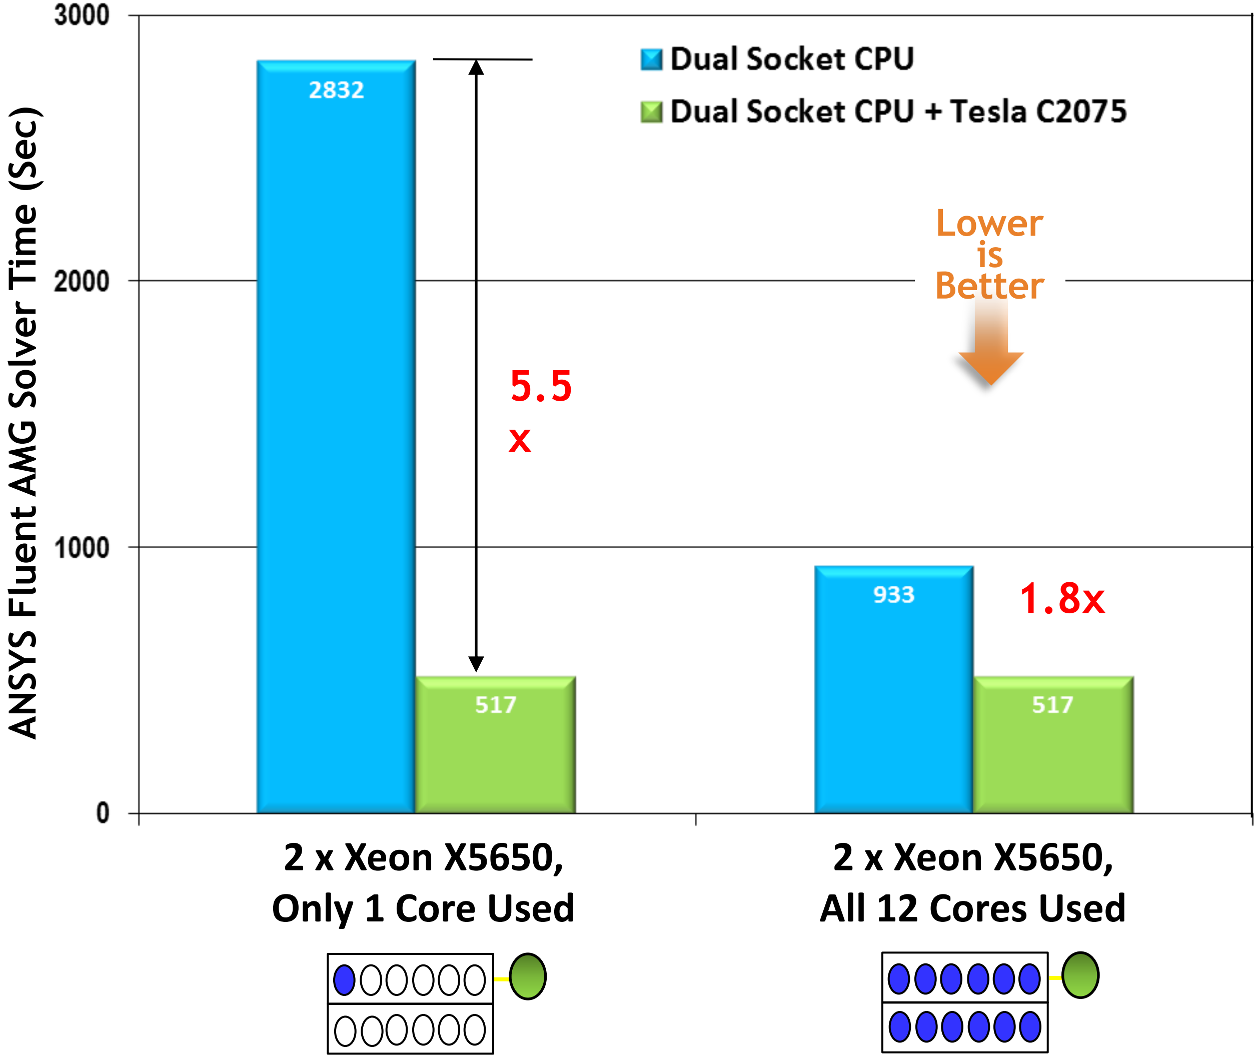

Helix Model

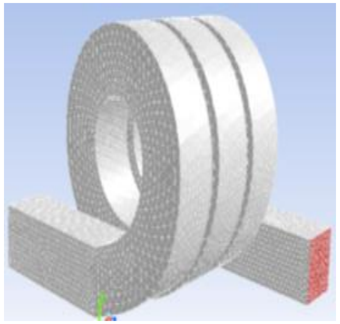

- Helix geometry
- 1.2M Tet cells
- Unsteady, laminar
- Coupled PBNS, DP
- AMG F-cycle on CPU
- AMG V-cycle on GPU

NOTE: All jobs solver time only

ANSYS Fluent Preview for 2 x CPU + 2 x Tesla K20X

ANSYS Fluent 15.0 Preview Performance – Results by NVIDIA, Feb 2013

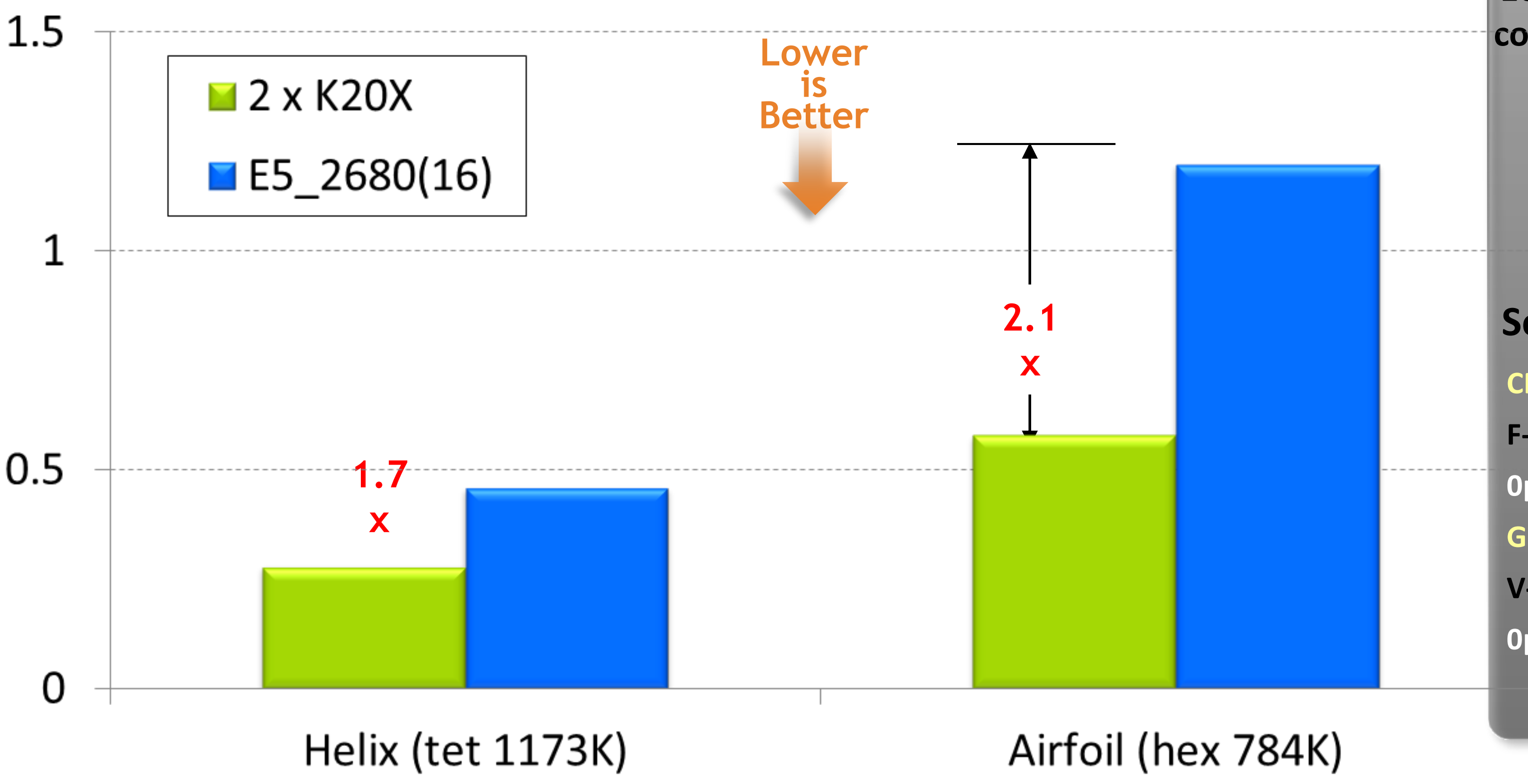

2 x E5_2680 SB CPUs, 16 cores total, only 2 cores used with GPUs

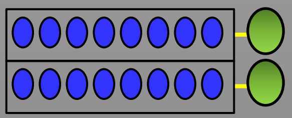

Solver settings:

CPU Fluent solver:

F-cycle, agg8, DILU, 0pre, 3post

GPU nvAMG solver:

V-cycle, agg8, MC-DILU, 0pre, 3post

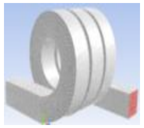

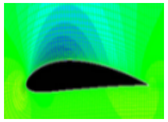

NOTE: Times for solver only

ANSYS Fluent Preview Scaling for 4 x Tesla K20X

ANSYS Fluent 15.0 Preview Performance – Results by NVIDIA, Mar 2013

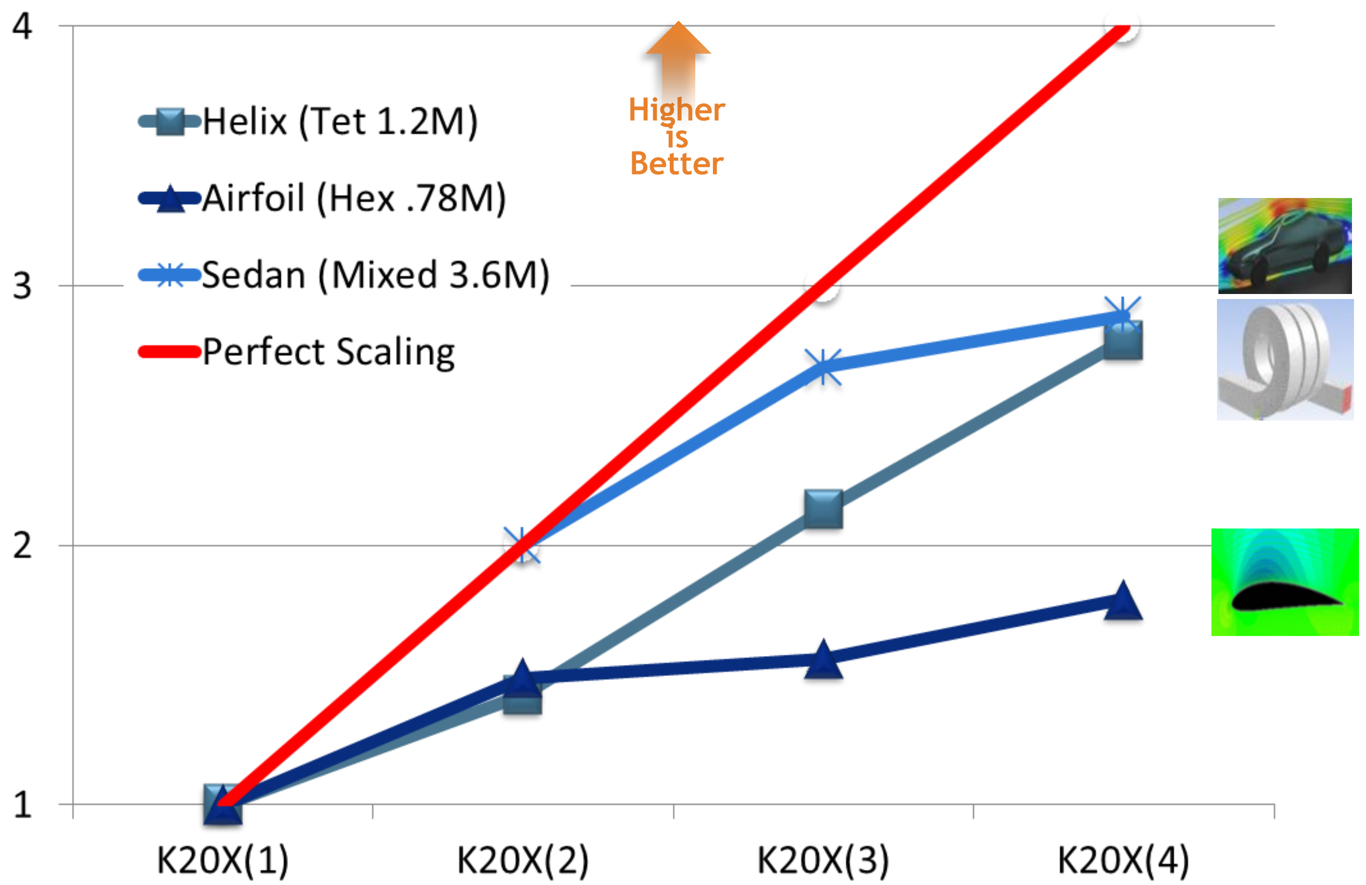

Hardware Setup:

- 2 server nodes
- 2 GPUs each node
- Infiniband network

GPU Solver Settings:

V-cycle, agg8/2,

MC-DILU, 0pre, 3post

NOTES:

- Results for solver only
- Sedan case starts with 2 GPUs

Hardware for TSNE Test

- **NVIDIA Tesla K40c**
  - Power : 245 W
  - Memory : 12 Gb
  - GPU Clock : 745 MHz

- **Workstation**
  - CPU : E5-2640W(2.50 Ghz)
  - Core : 12 (6 x 2)
  - Ram : 128Gb
  - Video : Standards VGA (on board)
  - OS : Windows 7 64bit
  - CUDA : 5.0

FLUENT Case 1

Exterior Flow Around a Passenger Sedan (FL5L2) - sedan_4m	
This benchmark represents the computation of the exterior flow field around a simplified model of a passenger sedan. The simulation geometry was used for the Japan External Aerodynamics competition. A viscous-hybrid grid with prismatic cells is used to adequately model the boundary layer regions.	
Number of cells	**3,618,080**
Cell type	hybrid
Models	k-epsilon turbulence
Solver	Coupled

ANSYS Fluent 15.0 Scaling for 4 x Tesla K40C

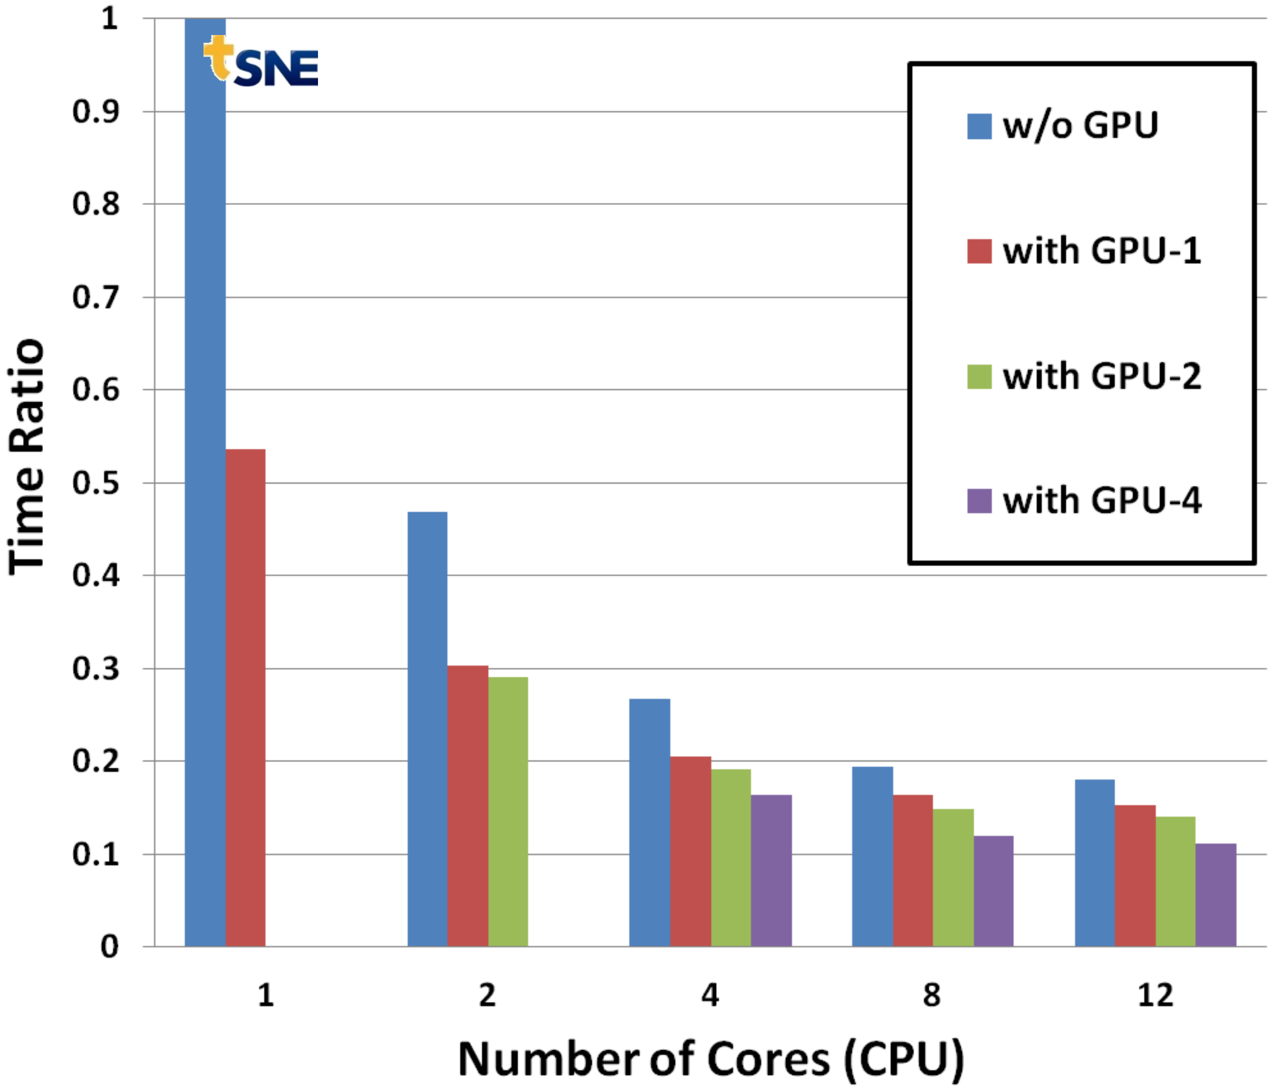

ANSYS FLUENT 15.0 with and without GPUs

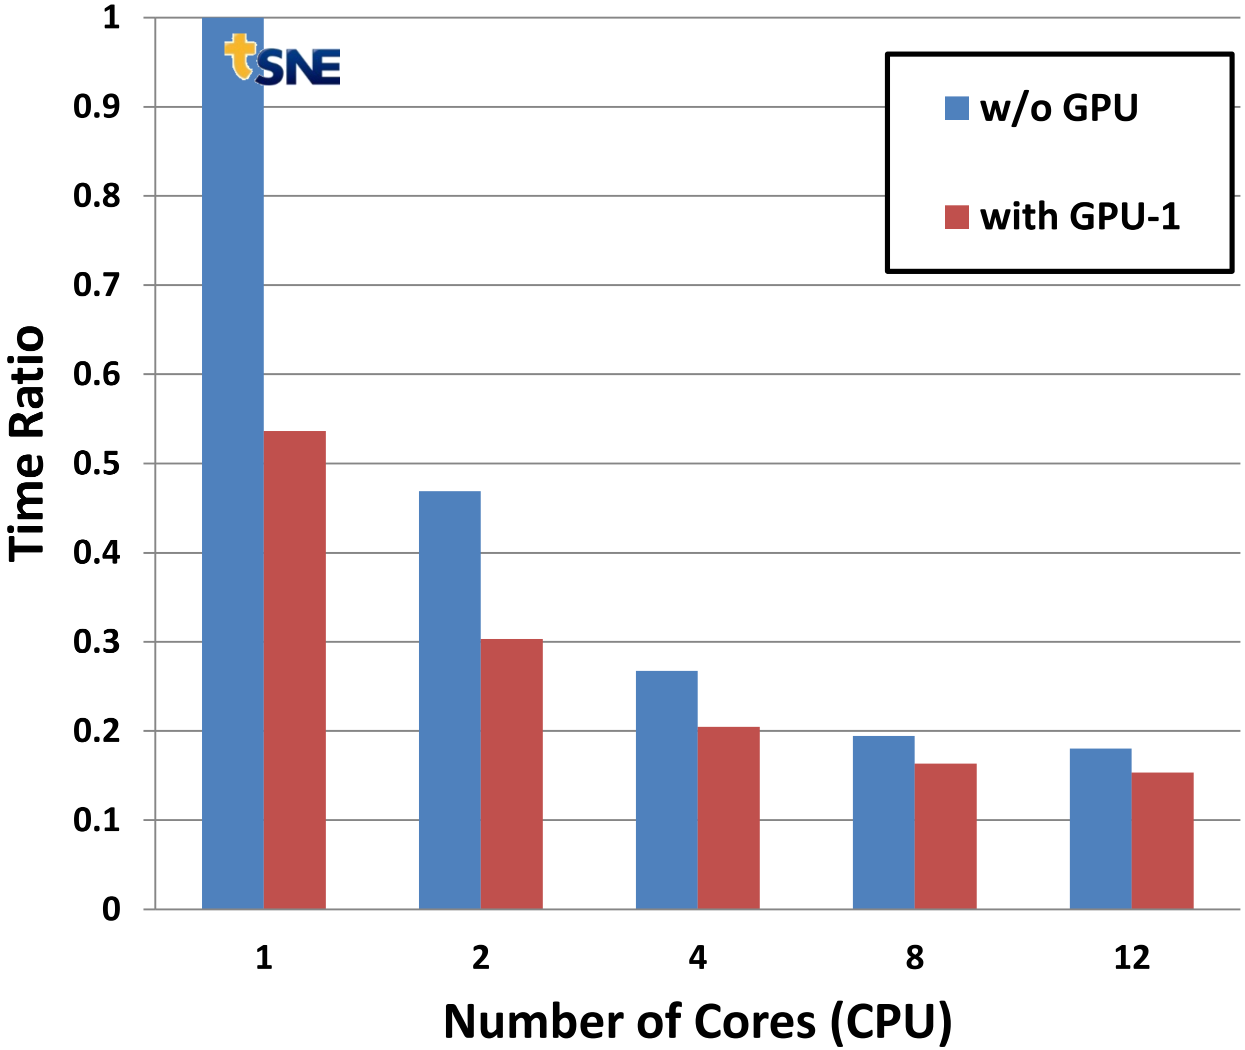

12 Parallel + 4 GPUs

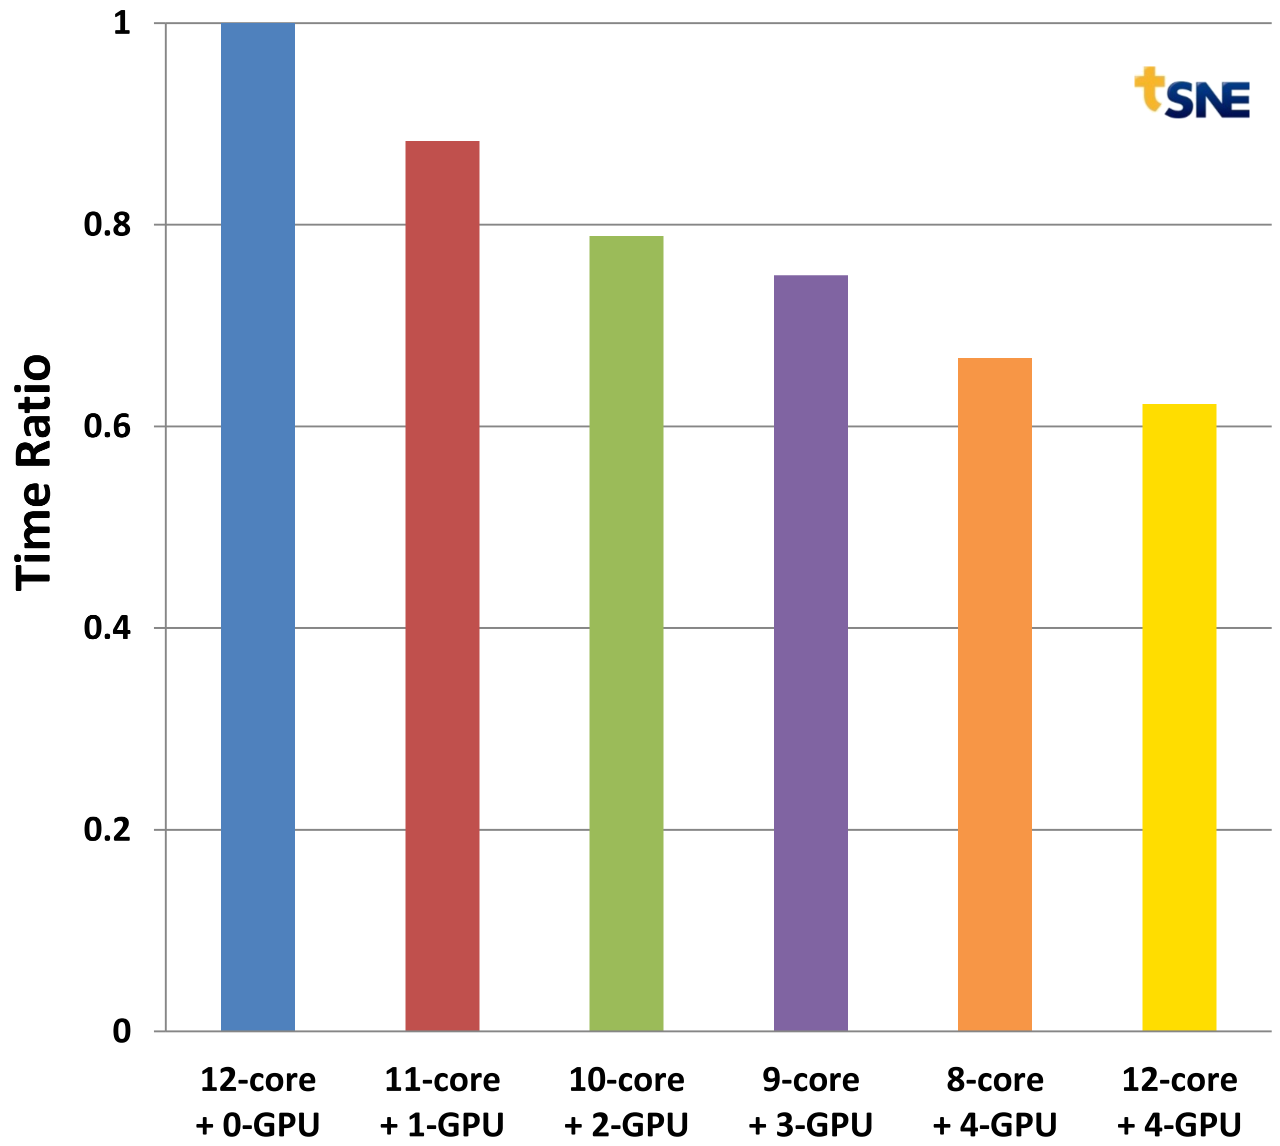

CPU vs GPUs

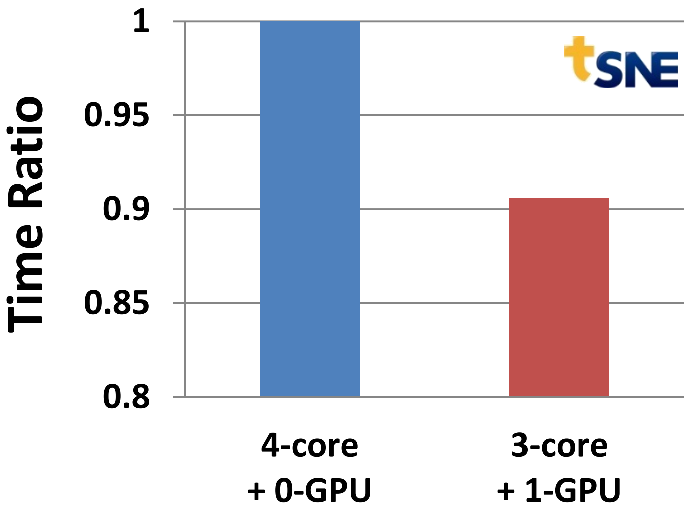
Time Ratio
1
0.95
0.9
0.85
0.8
4-core + 0-GPU
3-core + 1-GPU

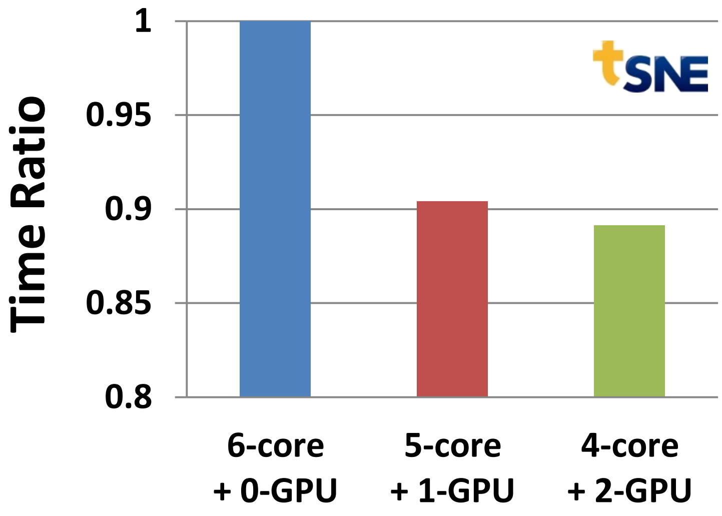
Time Ratio
1
0.95
0.9
0.85
0.8
6-core + 0-GPU
5-core + 1-GPU
4-core + 2-GPU

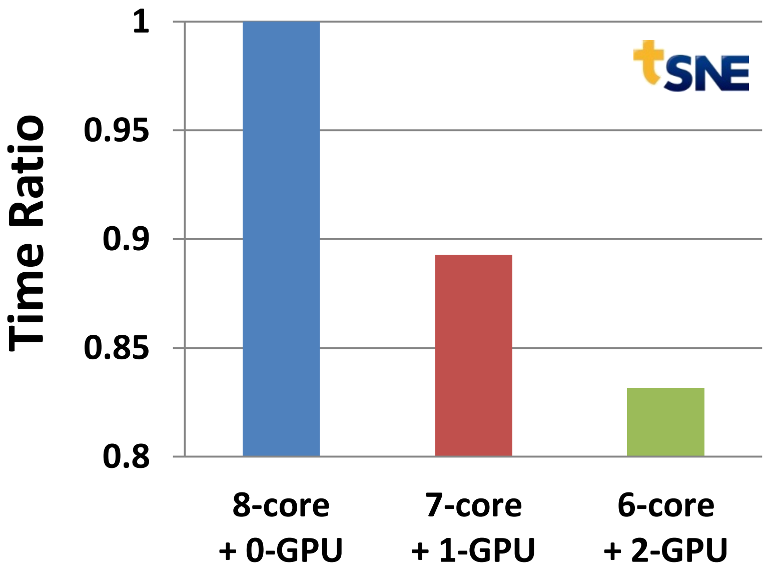
Time Ratio
1
0.95
0.9
0.85
0.8
8-core + 0-GPU
7-core + 1-GPU
6-core + 2-GPU

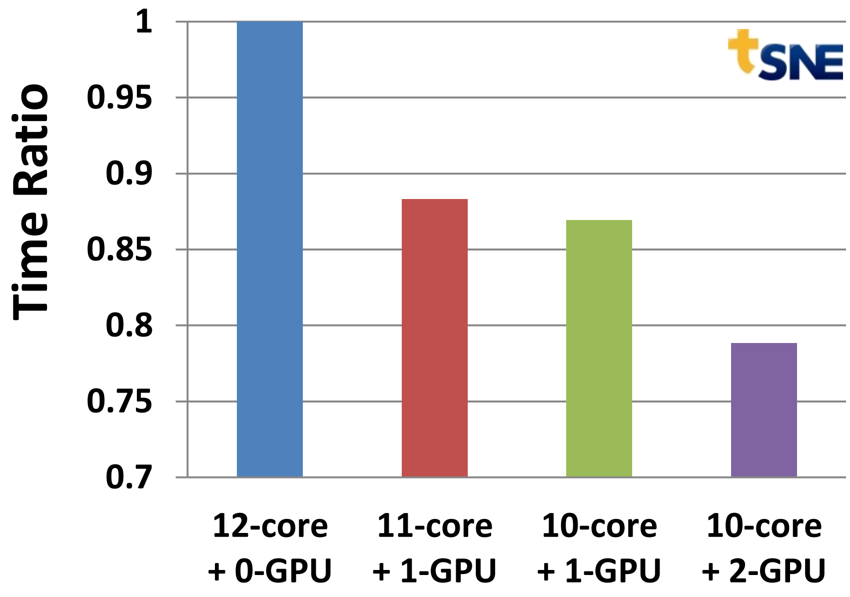
Time Ratio
1
0.95
0.9
0.85
0.8
0.75
0.7
12-core + 0-GPU
11-core + 1-GPU
10-core + 1-GPU
10-core + 2-GPU

Test Results

- Linux 와 Windows 에서 GPUs 를 사용하여 성능 향상을 보임
  - 12 parallel 사용시 CPU 만 사용한 경우 (core 12) 보다 GPUs 를 함께 사용한 경우(core 8 + 4 GPUs) 성능이 최대 60% 향상됨
  - 1 H/W 에서 multi-GPUs(4)를 추가 사용시 성능이 70% 향상됨
    : 12 core vs 12 core + 4 GPUs

- GPUs 사용 효과
  - 4 core 보다 3 core + 1 GPUs 가 빠름
  - 6 core 보다 5 core + 1 GPUs, 4 core + 2GPUs 가 빠름
  - 8 core 보다 7 core + 1 GPUs, 6 core + 2GPUs 가 빠름
  - 12 core 보다 11 core + 1 GPUs, 10 core + 1 or 2 GPUs 가 빠름

FLUENT Case 2

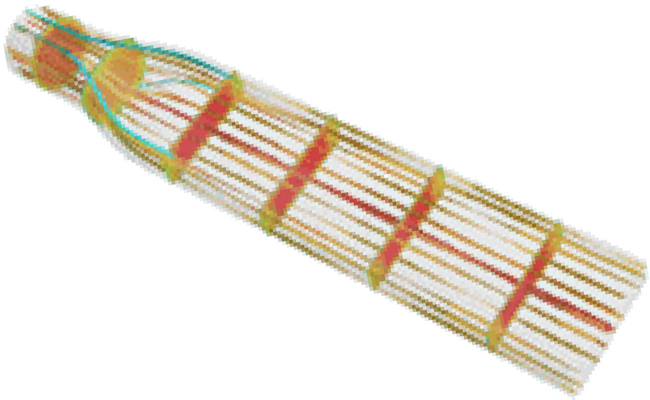

Turbulent Flow Through a Transition Duct (FL5L3)	
Turbulent flow of air through a duct is computed for this benchmark. The cross-sectional planes of the duct transition from a circle at the inlet to a rectangle at the outflow boundary. The	
Number of cells	**9,792,512**
Cell type	Hexahedral
Models	k-epsilon turbulence
Solver	Coupled

ANSYS Fluent 15.0 Scaling for 4 x Tesla K40C

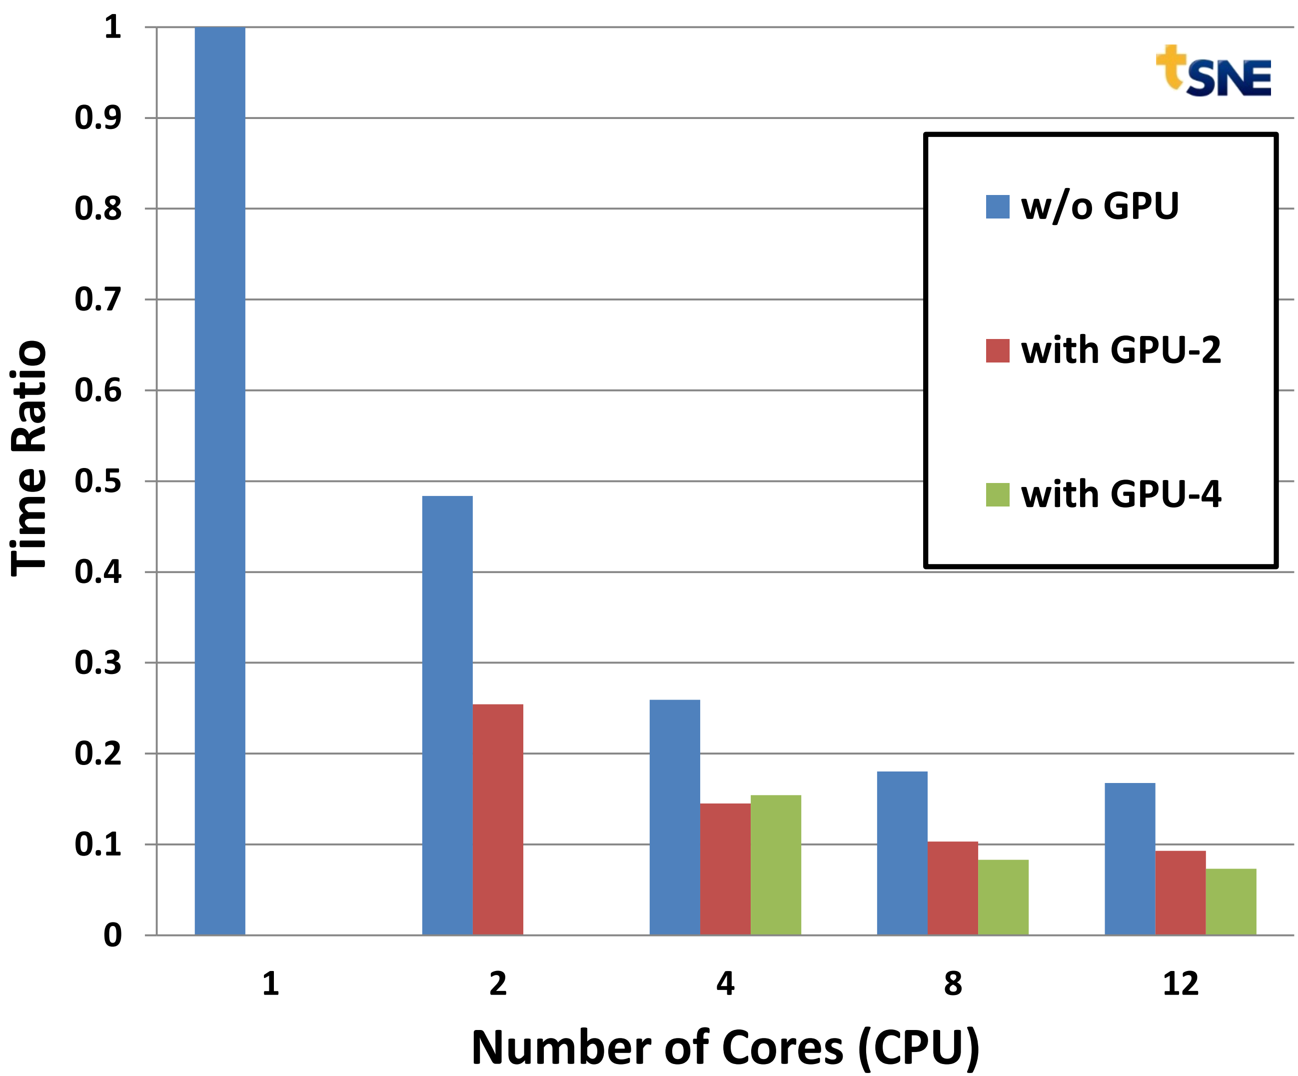

12 Parallel + 4 GPUs

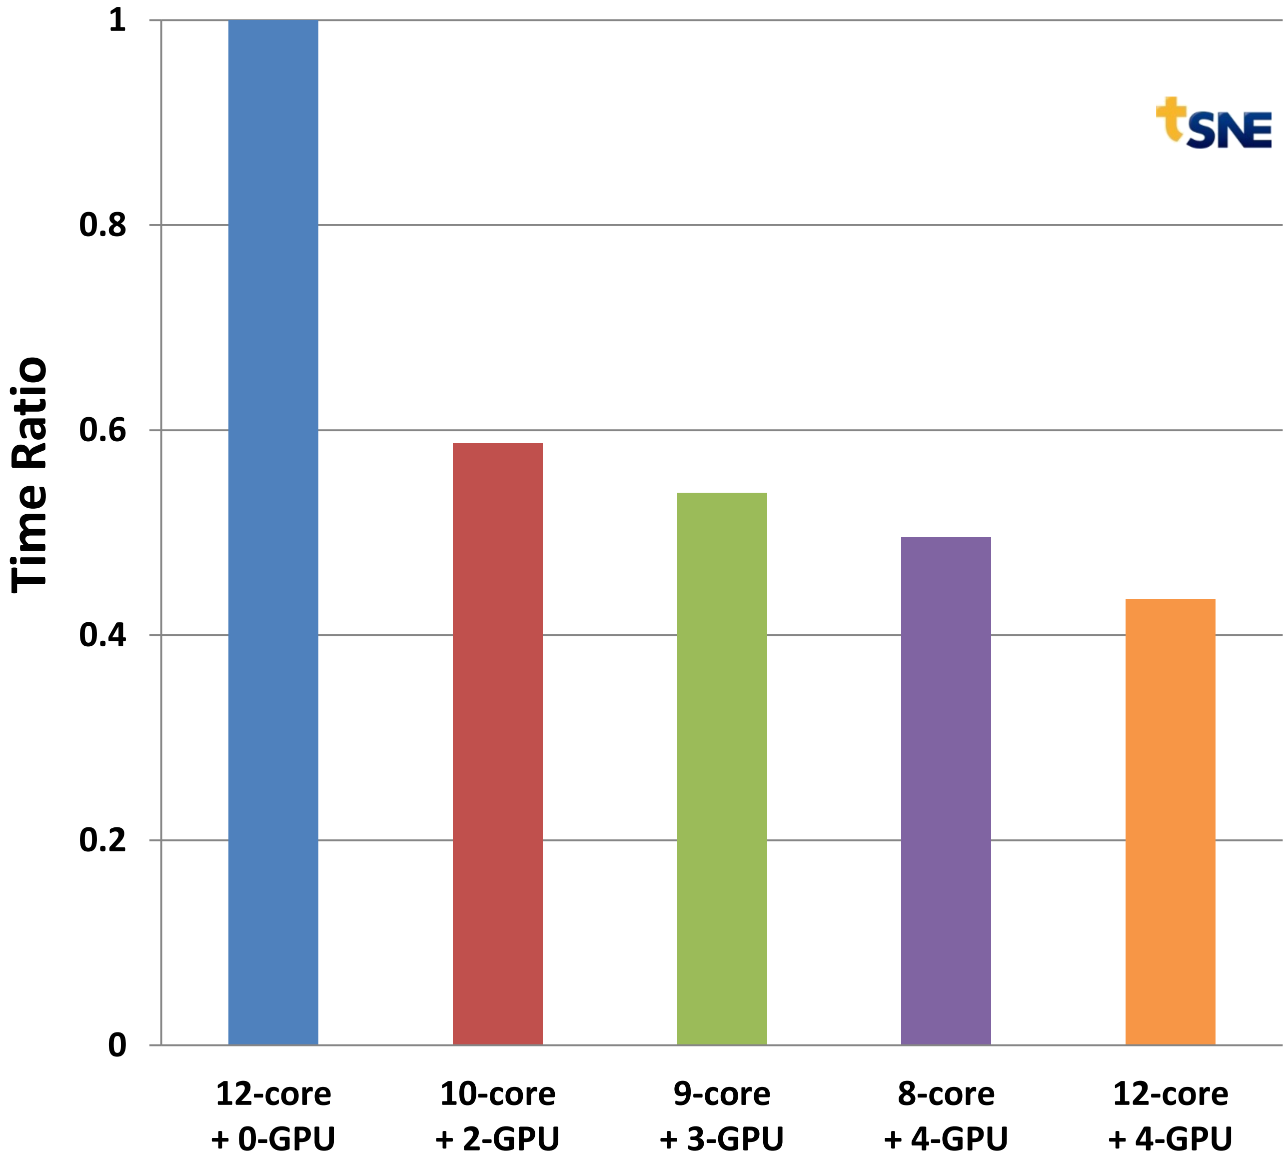

CPU vs GPUs

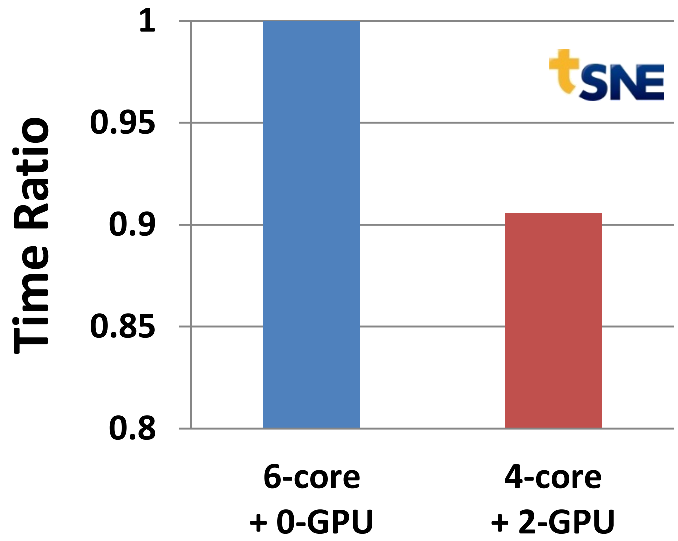

tSNE
Time Ratio
1
0.95
0.9
0.85
0.8
6-core + 0-GPU
4-core + 2-GPU

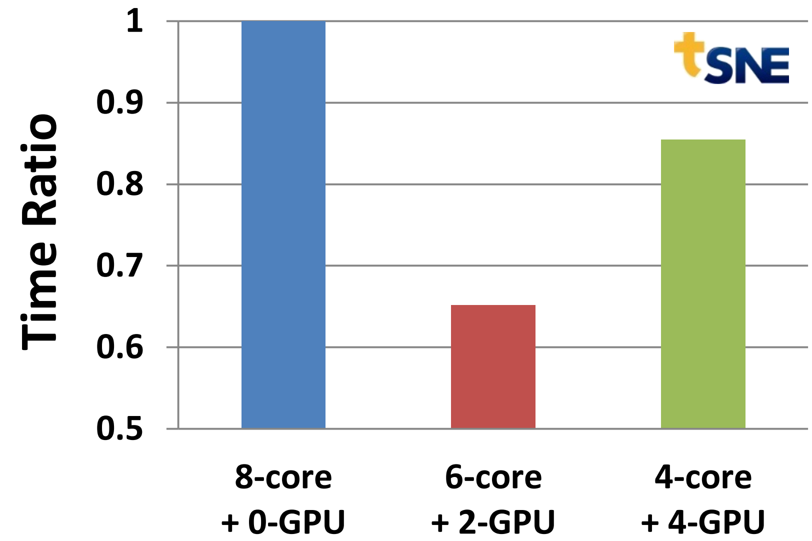

tSNE
Time Ratio
1
0.9
0.8
0.7
0.6
0.5
8-core + 0-GPU
6-core + 2-GPU
4-core + 4-GPU

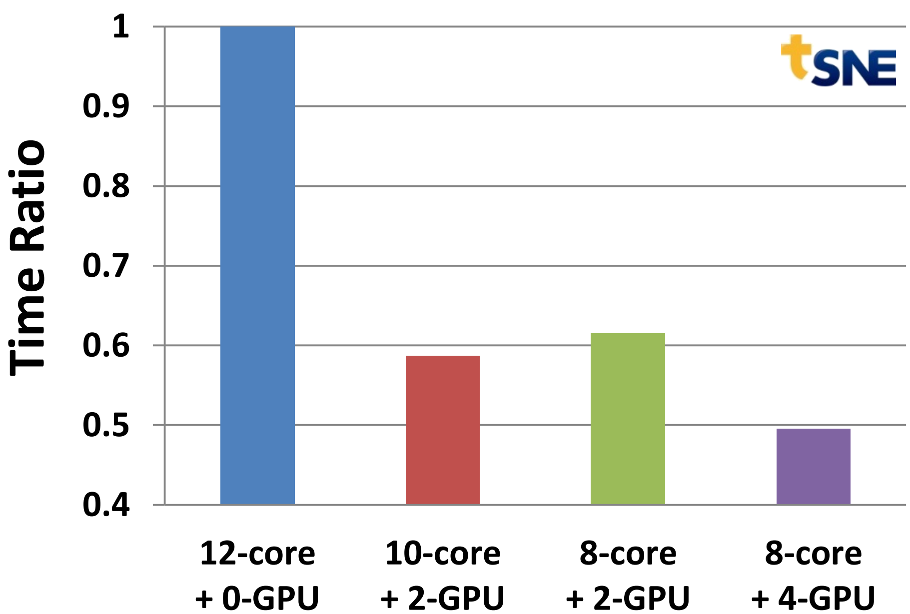

tSNE
Time Ratio
1
0.9
0.8
0.7
0.6
0.5
0.4
12-core + 0-GPU
10-core + 2-GPU
8-core + 2-GPU
8-core + 4-GPU

Test Results

- **GPUs 메모리가 12Gb 로 1 GPUs 로는 10 Million cells 해석이 어려움**
  - **Multi-GPU 를 이용하여 해석 가능**
- **Linux 와 Windows 에서 GPUs 를 사용하여 성능 향상을 보임**
  - **12 parallel 사용시 CPU 만 사용한 경우 (12 core) 보다 GPUs 를 함께 사용한 경우(8 core + 4 GPUs) 성능이 최대 100% 향상됨(2.0x)**
  - **1 H/W 에서 multi-GPUs(4)를 추가 사용시 성능이 130% 향상됨 (2.3x)**

    : 12 core vs 12 core + 4 GPUs
- **GPUs 사용 효과**
  - **6 core 보다 4 core + 2 GPUs 가 빠름**
  - **8 core 보다 6 core + 2 GPUs 가 빠름**
  - **12 core 보다 10 core + 2 GPUs, 8 core + 4 GPUs 가 빠름**

Conclusion

- **ANSYS CFD 15.0 병렬연산 성능이 향상 됨**

  - FLUENT (10,000 cells/core), CFX (50,000 nodes/ core)

- **GPUs를 사용하여 성능 향상을 보임**

  - **1 H/W** 에서 **multi-GPUs(4)**를 추가 사용시 성능이 **70~130%** 향상됨

  : 12 core vs 12 core + 4 GPUs

  - 한정된 라이선스 자원 내에서 최적의 계산 효율을 찾기 위해 **CPU + GPUs**에 대한 다양한 조합으로 테스트를 진행

  - **1 HPC Pack**을 보유한 경우 **8 CPU** 해석 보다는 **7 CPU + 1 GPUs** 또는 **6 CPU + 2 GPUs** 해석을 사용하는 것이 효율적임

지속적인 개발로 보다 좋은 성능과 안정성을 기대!!

THANK YOU

TAESUNG S&E,INC.